

THE TIME DOMAIN SPECTROSCOPIC SURVEY: VARIABLE OBJECT SELECTION AND ANTICIPATED RESULTS

ERIC MORGANSON¹, PAUL J. GREEN¹, SCOTT F. ANDERSON², JOHN J. RUAN², ADAM D. MYERS³, MICHAEL ERACLEOUS⁴, BRANDON KELLY⁵, CARLOS BADENES⁶, EDUARDO BAÑADOS⁷, MICHAEL R. BLANTON⁸, MATTHEW A. BERSHADY⁹, JURA BORISOVA¹⁰, WILLIAM NIELSEN BRANDT⁴, WILLIAM S. BURGETT¹¹, KENNETH CHAMBERS¹², PETER W. DRAPER¹³, JAMES R. A. DAVENPORT², HEATHER FLEWELLING¹², PETER GARNAVICH¹⁴, SUZANNE L. HAWLEY², KLAUS W. HODAPP¹², JEDIDAH C. ISLER^{15, 1}, NICK KAISER¹¹, KAREN KINEMUCHI¹⁶, ROLF P. KUDRITZKI¹², NIGEL METCALFE¹², JEFFREY S. MORGAN¹¹, ISABELLE PÂRIS¹⁷, MAHMOUD PARVIZI¹⁸, RADOSŁAW POLESKI¹⁹, PAUL A. PRICE²⁰, MARA SALVATO²¹, TOM SHANKS¹³, EDDIE F. SCHLAFLY⁷, DONALD P. SCHNEIDER⁴, YUE SHEN^{22, 23}, KEIVAN STASSUN¹⁸, JOHN T. TONRY¹², FABIAN WALTER⁷, AND CHRIS Z. WATERS¹²

Accepted by the Astrophysical Journal

ABSTRACT

We present the selection algorithm and anticipated results for the Time Domain Spectroscopic Survey (TDSS). TDSS is an SDSS-IV eBOSS subproject that will provide initial identification spectra of approximately 220,000 luminosity-variable objects (variable stars and AGN) across 7,500 deg² selected from a combination of SDSS and multi-epoch Pan-STARRS1 photometry. TDSS will be the largest spectroscopic survey to explicitly target variable objects, avoiding pre-selection on the basis of colors or detailed modeling of specific variability characteristics. Kernel Density Estimate (KDE) analysis of our target population performed on SDSS Stripe 82 data suggests our target sample will be 95% pure (meaning 95% of objects we select have genuine luminosity variability of a few magnitudes or more). Our final spectroscopic sample will contain roughly 135,000 quasars and 85,000 stellar variables, approximately 4,000 of which will be RR Lyrae stars which may be used as outer Milky Way probes. The variability-selected quasar population has a smoother redshift distribution than a color-selected sample, and variability measurements similar to those we develop here may be used to make more uniform quasar samples in large surveys. The stellar variable targets are distributed fairly uniformly across color space, indicating that TDSS will obtain spectra for a wide variety of stellar variables including pulsating variables, stars with significant chromospheric activity, cataclysmic variables and eclipsing binaries. TDSS will serve as a pathfinder mission to identify and characterize the multitude of variable objects that will be detected photometrically in even larger variability surveys such as LSST.

Subject headings:

emorganson@cfa.harvard.edu

¹ Harvard Smithsonian Center for Astrophysics, 60 Garden St, Cambridge, MA 02138, USA

² Department of Astronomy, University of Washington, Box 351580, Seattle, WA 98195, USA

³ Department of Physics and Astronomy, University of Wyoming, Laramie, WY 82071, USA

⁴ Department of Astronomy & Astrophysics, 525 Davey Laboratory, The Pennsylvania State University, University Park, PA 16802, USA

⁵ Department of Physics, Broida Hall, University of California, Santa Barbara, CA 93106-9530, USA

⁶ Department of Physics and Astronomy and Pittsburgh Particle Physics, Astrophysics and Cosmology Center (PITT PACC), University of Pittsburgh, 3941 O'Hara St, Pittsburgh, PA 15260, USA

⁷ Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany

⁸ Center for Cosmology and Particle Physics, Department of Physics, New York University, 4 Washington Place, New York, NY 10003, USA

⁹ Department of Astronomy, University of Wisconsin, 475 N. Charter St., Madison, WI 53706, USA

¹⁰ Instituto de Física y Astronomía, Universidad de Valparaíso, Av. Gran Bretaña 1111, Playa Ancha, Casilla 5030, and Millennium Institute of Astrophysics (MAS), Santiago, Chile

¹¹ GMTO Corp, Suite 300, 251 S. Lake Ave, Pasadena, CA 91101, USA

¹² Institute for Astronomy, University of Hawaii at Manoa, Honolulu, HI 96822, USA

¹³ Department of Physics, University of Durham Science Laboratories, South Road Durham DH1 3LE, UK

¹⁴ Department of Physics, University of Notre Dame, Notre Dame, IN 46556, USA

¹⁵ Syracuse University, Syracuse, NY 13244, USA

¹⁶ Apache Point Observatory, P.O. Box 59, Sunspot, NM 88349, USA

¹⁷ INAF - Osservatorio Astronomico di Trieste, Via G. B. Tiepolo 11, I-34131 Trieste, Italy

¹⁸ Department of Physics & Astronomy, Vanderbilt University, VU Sta-

tion B 1807, Nashville, TN, USA

¹⁹ Department of Astronomy, Ohio State University, 140 West 18th Avenue, Columbus, OH 43210, USA

²⁰ Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

²¹ Max Planck institute for extraterrestrial Physics, Giessenbachstr. 1, Garching, D-85748, Germany

²² Carnegie Observatories, 813 Santa Barbara Street, Pasadena, CA 91101, USA

²³ Kavli Institute for Astronomy and Astrophysics, Peking University, Beijing 100871, China

1. INTRODUCTION

Variability in optical luminosity is an important behavior in many astronomical objects, and enhancing our understanding of the physics of a variety of systems. In this paper, we discuss objects whose optical luminosity varies by a tenth of a magnitude or more on time scales of a year or less, a level that can be easily measured with ground-based observations, and refer to them as "variable objects". This term encompasses both stellar variables and AGN. The large majority of AGN (especially quasars) and approximately one percent of stars satisfy this definition. Quasars and other AGN generally vary stochastically in optical bands by up to several tenths of a magnitude over months and years (Giveon et al. 1999; Vanden Berk et al. 2004). The main cause of quasar variability in the optical continuum is instability in the accretion disk (Rees 1984; Kawaguchi et al. 1998; Pereyra et al. 2006; Ruan et al. 2014). Blazars, generally accepted to be AGN whose relativistic jets point along the line of sight (Antonucci 1993; Urry & Padovani 1995), vary due to Doppler beaming of their jet emission (Ulrich et al. 1997). Microlensing by stars in intervening lensing galaxies (Wambsganss 2006; Morgan et al. 2010) can also contribute to AGN variability in some cases.

Stellar variability is produced by a large variety of physical processes. Chromospheric magnetic fields cause flaring stellar activity (Schatzman 1962; Wilson 1963; Baliunas et al. 1995; Hall et al. 2009; Mathur et al. 2014) that produces significant optical variability, particularly in younger late-type stars. Periodically pulsating variable stars exhibit large amplitude variability caused by the κ mechanism in which a star's atmospheric opacity varies periodically (Zhevakin 1959). These are more likely to appear as early-type stars, and the most famous pulsators, RR Lyrae and Cepheid variables, are commonly used as "standard candle" distance probes (Hubble 1929; Rodgers 1957; Pritchett & van den Bergh 1987; Smith 1995; Freedman et al. 2001; Sesar et al. 2010). Cataclysmic variables (CVs) are binaries in which a white dwarf accretes material from its companion producing occasional outbursts that can generate several magnitudes of variability (Mumford 1963; Connon Smith 2007; Knigge 2011). CV donor stars can appear as a wide variety of stellar types although most CVs involve a red dwarf or giant. Eclipsing binaries can also produce significant periodic variability (Stephenson 1960; Debosscher et al. 2011; Beck et al. 2014) across all stellar types.

Because of its astrophysical importance, variability has become the focus of many recent and upcoming photometric surveys in which the same region of sky is imaged multiple times. A series of small (20-100 deg²) surveys including the Faint Sky Variability Survey (Groot et al. 2003) and MACHO (Alcock et al. 2001) have obtained hundreds to thousands of photometric measurement epochs. The Kepler Mission (Borucki et al. 2010) is probing similarly sized areas with much greater photometric precision and tens of thousands of observation epochs. OGLE I-OGLE IV (Udalski et al. 2008; Wyrzykowski et al. 2014), the QUEST RR-Lyrae Survey (Vivas et al. 2004), the Sloan Digital Sky Survey (SDSS, York et al. 2000) Stripe 82 (Sesar et al. 2007) and the VISTA Variables in the Vía Láctea ESO Public Survey (Catelan et al. 2011) cover 2,000, 700, 290 and 560 deg² respectively with each providing of order 100 measurement epochs per source. Recently, a number of "full sky" (at least 10,000 deg²) variability surveys have been completed. ROTSE-I (Akerlof et al. 2000; Woźniak et al. 2004b), The La Silla-QUEST Variability

Survey in the Southern Hemisphere (Hadjiyska et al. 2012), the Catalina Sky Survey (Drake et al. 2009), the Palomar Transient Factory (PTF, Law et al. 2009), All-Sky Automated Survey (ASAS, Pojmanski 2002), the Lincoln Near-Earth Asteroid Research survey (LINEAR, Palaversa et al. 2013) and Pan-STARRS1 (PS1, Kaiser et al. 2002, 2010) obtain between 50 and 400 measurements per object. Of these, PS1 is the deepest and covers the largest area, and PS1 data will be the focus of this paper. In the near future, the Gaia mission (Lindegren et al. 2008) and the Large Synoptic Survey Telescope (LSST Science Collaboration et al. 2009) will extend full sky surveys to greater precision, more rapid cadences and much fainter limits.

These photometric surveys have been accompanied by many large spectroscopic surveys. The SDSS-III Baryon Oscillation Spectroscopic Survey (BOSS, Dawson et al. 2013), its SDSS-IV extension eBOSS (eBOSS, Dawson et al. 2015, *in preparation*) and the LAMOST ExtraGalactic Surveys (LEGAS, Wang et al. 2009) will eventually take 1.3×10^6 spectra of quasars, which are generally variable. SDSS has also taken 2.4×10^5 optical stellar spectra in the Sloan Extension for Galactic Understanding and Exploration (SEGUE, Yanny et al. 2009) and will take 10^5 high resolution infrared spectra with the APO Galactic Evolution Experiment (APOGEE, Zasowski et al. 2013). The Bulge Radial Velocity Assay (BRAVA, Kunder et al. 2012), the Radial Velocity Experiment (RAVE, Kordopatis et al. 2013), the LAMOST Experiment for Galactic Understanding and Exploration (LEGUE, Deng et al. 2012) and the GALactic Archaeology with HERMES survey (GALAH, Zucker et al. 2012) will obtain between 10^4 and 2.5×10^6 stellar spectra each. We expect roughly 1% of the stars in each of these surveys to satisfy our definition of variable. Finally, the Gaia mission will obtain high resolution ($R \approx 11,500$) narrow filter ($8,470 \text{ \AA} < \lambda < 8,740$) and low resolution ($10 < R < 200$) broad filter ($3,300 \text{ \AA} < \lambda < 10,000$) spectroscopy of $10^8 V < 17$ objects. These spectra will provide precise radial velocity measurements and generally characterize a wide variety of astrophysical objects, but for the broad variety of galactic and extragalactic variable objects we target may be less useful at characterizing e.g., specific absorption and emission lines that fall outside the narrow high resolution spectra.

Despite these dedicated photometric variability surveys and similarly large spectroscopic surveys, large spectroscopic surveys of variable objects are somewhat lacking. There have been variability-selected samples of quasars (e.g. Palanque-Delabrouille et al. 2011a) and RR Lyrae stars (e.g. Drake et al. 2013) as well as relatively small SDSS spectroscopic variability studies of subdwarfs (Geier et al. 2011), white dwarf main-sequence binaries (Rebassa-Mansergas et al. 2011), white dwarfs (Badenes et al. 2009; Mullally et al. 2009; Badenes et al. 2013) and field stars more generally (Pourbaix et al. 2005). But these surveys have been relatively small in size and have used color information, spectra or specific light curve character to target specific types of variables.

The Time Domain Spectroscopic Survey (TDSS) has been designed to widen the scope of spectroscopic surveys of variable objects and will soon become the largest medium resolution ($R \approx 2,000$), broad wavelength ($3,600 \text{ \AA} < \lambda < 10,400 \text{ \AA}$) spectroscopic survey of variable objects. This survey, a sub-project of the SDSS-IV Extended Baryon Oscillation Spectroscopic Survey (eBOSS), will cover 7,500 deg² and include 220,000 variability-selected targets with no focus on any spe-

cific variability or photometric type in target selection. TDSS is not well-suited to spectroscopic identification of rapid transients, because plug plates to accommodate the 1,000 spectroscopic fibers must be drilled well in advance of observations.

Roughly 90% of TDSS targets will be TDSS’s Single Epoch Spectroscopy (SES) targets, for which TDSS will produce a single discovery (identification/classification) spectrum. For bright, quickly varying targets within this sample, we will also be able to study spectroscopic variability by examining the spectroscopic sub-exposures taken over hours or sometimes several nights. This TDSS SES sample is designed to be a probe of general optical variability and will be the subject of this paper.

The remaining 10% of TDSS spectra will be drawn from one of TDSS’s nine Few Epoch Spectroscopy (FES) projects, a series of smaller ($\approx 1,000$ targets each) samples with previous SDSS spectroscopy for which TDSS will obtain another spectrum for two and occasionally three-epoch comparison. The FES projects are each designed to probe a specific type of variable object and science topic. These nine current projects are devoted to:

- Radial velocity variation in dwarf carbon stars
- M-dwarf white dwarf binaries
- Activity in ultracool dwarfs on decadal time scales
- Stars with more than 0.2 magnitudes of variability
- Broad absorption line trough variability (as in Filiz Ak et al. 2013) in quasars
- Balmer line variability in high signal to noise quasars
- Double-peaked broad emission line quasars
- Searching for binary black hole quasars via Mg II line velocity shifts
- Quasars with more than 0.7 magnitudes of variability

The details of these FES projects will be addressed in future papers.

In this paper, we describe how the TDSS SES Project (subsequently referred to as simply TDSS) produces a large sample of photometric variable objects with a broad range of variability types while avoiding spurious, non-astrophysical “variability” in its target selection. In Section 2 we outline TDSS’s role in eBOSS, the larger SDSS-IV optical spectroscopy project. In Section 3 we demonstrate how the combination of SDSS and Pan-STARRS1 photometry allows the construction of a $7,500 \text{ deg}^2$, relatively uniform sample, and we describe our algorithm for quantifying variability into a single metric in Section 4. In Section 5 we present our ultimate target prioritization. We estimate our survey purity (fraction of candidates that genuinely vary by a few tenths of magnitudes) and show how it varies across the sky in Section 6. We describe the selection of a small subsample of i -band dropouts that would have been missed by our algorithm without special effort in Section 7. We statistically classify our complete list of targets by their colors in Section 8 and discuss how our selection percentage varies as a function of color in Section 9. Finally, we compare the targets selected by our algorithm using our dataset to small sets of known variable objects and objects with existing SDSS spectra from SDSS Stripe 82 in Section 10.

2. TDSS AND EBOSS

TDSS is a subprogram of the Extended Baryon Oscillation Spectroscopic Survey (eBOSS). eBOSS is an SDSS-IV project designed to perform a variety of cosmological measurements with spectroscopy of quasars (Myers et al. 2015, *in preparation*), luminous red galaxies (Prakash et al. 2015, *in preparation*), X-ray emitting quasars and cluster galaxies (Menzel et al. 2015 *in preparation*, Finoguenov et al. 2015, *in preparation* and Clerc et al. 2015, *in preparation*) and emission line galaxies (Comparat et al. 2013). TDSS will be paired with the main eBOSS survey (shown in Fig. 1) and is planned to cover a total of $7,500 \text{ deg}^2$ in the Northern and Southern Galactic Caps. eBOSS devotes 10 fibers deg^{-2} to TDSS-only targets. But TDSS also selects an additional 23 TDSS-joint targets deg^{-2} that have previous SDSS spectroscopy or are part of the main eBOSS quasar target list most of which is selected using colors alone with the $XDQSOz$ algorithm (Bovy et al. 2012). A small number of eBOSS quasars are also selected using a combination of colors and optical variability from the Palomar Transient Factory (Palanque-Delabrouille et al. 2011b). The full TDSS sample will thus include 33 objects deg^{-2} . See Section 6 for more details.

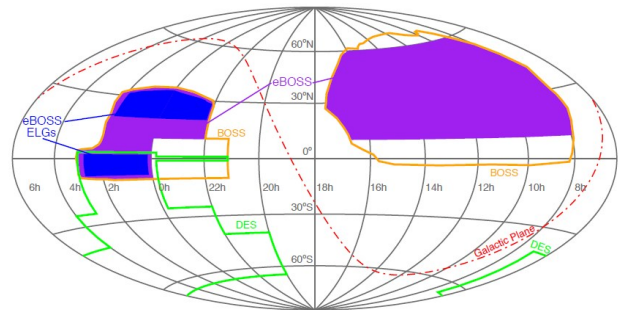


FIG. 1.— The planned eBOSS (and by extension TDSS) area is shown in blue and purple. The blue area may be sampled twice for the eBOSS Emission Line Galaxy (ELG) project. The eBOSS predecessor, BOSS, is outlined in orange, and the Dark Energy Survey, which may be of interest in ELG targeting, is outlined in green.

Spectroscopy for the main TDSS sample will be obtained as part of the eBOSS schedule on the BOSS spectrograph (Smee et al. 2013). At the TDSS $i = 21$ magnitude limit, we will obtain per pixel signal to noise ratios of 5 or better (Dawson et al. 2013) (typical pixel size is roughly 1 \AA). We use an $i = 17$ bright limit to prevent saturation and signal leaking between adjacent fibers. The spectra cover $3,700 \text{ \AA} < \lambda < 10,400 \text{ \AA}$ in two channels (red and blue). The spectrograph’s resolution runs from $R=1,560$ at $3,700 \text{ \AA}$ to $R=2,270$ at $6,000 \text{ \AA}$ (blue channel), and from $R=1,850$ at $6,000 \text{ \AA}$ to $R=2,650$ at $9,000 \text{ \AA}$ (red channel). These spectra will be easily good enough measure continua, major absorption and emission features, quasar redshifts and stellar velocities (to better than 50 km s^{-1}).

3. THE SDSS-PS1 DATASET

In order to measure optical variability, TDSS uses a combination of single-epoch SDSS and multi-epoch PS1 photometry. We use SDSS photometry from SDSS Data Release 9 (Gunn et al. 1998; York et al. 2000; Gunn et al. 2006; Aihara et al. 2011; Eisenstein et al. 2011; Ahn et al. 2012). SDSS DR9 covers $14,555 \text{ square degrees}$ in the u, g, r, i and z filters which span the $3,000 \text{ \AA} < \lambda < 10,000 \text{ \AA}$ spectral range

Filter	SDSS	PS1 Exposure	PS1 Mean	PS1 Stack
<i>u</i>	21.2	—	—	—
<i>g</i>	22.3	21.2	21.7	22.4
<i>r</i>	21.8	21.0	21.7	22.2
<i>i</i>	21.4	20.8	21.5	22.0
<i>z</i>	19.9	20.1	20.7	21.3
<i>y</i>	—	19.1	19.7	20.3

TABLE 1

Median 10σ Limiting AB PSF Magnitudes of SDSS, PS1 3π single exposures, the PS1 3π mean catalog and the current PS1 3π stack. Similarly-named filters from different surveys are not exactly the same.

(Fukugita et al. 1996). The imaging footprint covers most of the high Galactic longitude area north of declination -10° . Throughout this paper, we use *u*, *g*, *r*, *i* and *z* to refer to the SDSS magnitudes and not the (very similar) PS1 analogs.

The PS1 3π survey (Kaiser et al. 2002, 2010; Chambers 2011) covers its 30,000 deg² area north of declination -30° . This region includes the entire SDSS survey imaging footprint. The PS1 g_{P1} , r_{P1} , i_{P1} and z_{P1} filters cover the $4000 \text{ \AA} < \lambda < 9200 \text{ \AA}$ spectral range similarly to the corresponding SDSS *g*, *r*, *i* and *z* filters. PS1 also has a y_{P1} filter which, including the spectral response of the camera, covers $9200 \text{ \AA} < \lambda < 10500 \text{ \AA}$. These PS1 filters are described in detail in Tonry et al. (2012). The PS1 survey takes four exposures per year for 3.5 years with each of the $g_{P1}r_{P1}i_{P1}z_{P1}y_{P1}$ filters (non-simultaneously) and fills approximately 90% of the 30,000 deg² area in each band. The missing area is mostly due to non-detection areas on the camera plane and weather restricting the survey to two or rarely zero exposures per filter in some areas of the sky. Individual PS1 exposures are generally shallower than analogous SDSS images. However, the 10σ limiting PSF magnitudes of the PS1 average catalogs, produced by taking a weighted average of individual detections rather than stacking the images (PS1 image stacking is still being developed), are well-matched to the SDSS single-exposure limits as summarized in Table 1.

In this work, we use an updated version of the "übercalibrated" PS1 data from Schlafly et al. (2012), which includes the PS1 data up through July 2013 (using PV1 of the PS1 pipeline) and is calibrated absolutely to 0.02 magnitudes or better. This database excludes detections flagged by PS1 as cosmic rays, edge effects and other defects.

To convert between SDSS and PS1 magnitudes, we use the conversions from Finkbeiner et al. (2014) which follow the equation

$$m_{P1} - m_{SDSS} = a_0 + a_1 gi + a_2 gi^2 + a_3 gi^3, \quad (1)$$

$$gi = g - i.$$

where $m = griz$ and a_{0123} are in Table 2. Tonry et al. (2012) also provide a similar conversion from SDSS to PS1 calculated from PS1 filter curves, but we use the Finkbeiner equations because they are optimized to be accurate for a broad stellar population, and because they are calculated within the Schlafly et al. (2012) übercalibrated system. For the non-varying stars for which these coefficients were fit, these conversions are accurate to 0.01 magnitudes or better. We add this 0.01 mag in quadrature to our statistical error. When comparing SDSS and PS1 magnitudes, we convert them to standard logarithmic magnitudes, rather than the default asinh-based "Luptitudes" that SDSS reports (Lupton et al. 1999).

All database analysis and cross-matching of surveys is per-

Filter	a_0	a_1	a_2	a_3
<i>g</i>	0.00128	-0.10699	0.00392	0.00152
<i>r</i>	-0.00518	-0.03561	0.02359	-0.00447
<i>i</i>	0.00585	-0.01287	0.00707	-0.00178
<i>z</i>	0.00144	0.07379	-0.03366	0.00765

TABLE 2

The coefficients used to convert from SDSS magnitude to PS1 magnitudes in Eq. 1. Ensemble error bars are insignificant; for individual stars, these conversions are good to 0.01 magnitudes.

formed with the Large Survey Database software (LSD Juric 2011). LSD is a versatile, parallelized, python-based database module optimized for astronomical querying and cross-matching. We compare PS1 and SDSS PSF magnitudes in all cases, only work with objects that are unresolved in SDSS (morphology type "star") and match PS1 and SDSS objects with a radius of 1.5".

Fig. 2 shows a typical SDSS-PS1 light curve for a non-variable and a variable object. The SDSS and PS1 magnitudes are consistent for the non-variable, confirming at least the approximate validity of the Finkbeiner conversions in this case. It is easily discernible that the variable object is varying in the PS1 data, but it is also clear that the very sparse PS1 sampling prevents a detailed characterization of a single object's variability (e.g. determining a period). This limitation, combined with a desire to avoid biasing our sample to any specific variability type, led to the relatively simple variability criteria described in Section 4.

3.1. PS1 Photometric Uncertainties

Accurate photometric uncertainties are very important for variability measurements. If we overestimate photometric error, we will underestimate variability and vice versa. To assess the level of spurious variability induced by incorrect error bars in PS1, we draw from the PS1 catalog a population of 4,032,258 (theoretically) constant photometric F stars which satisfy the following SDSS criteria (all magnitudes PSF magnitudes and are dereddened using the Schlegel et al. (1998) extinction map):

$$16 < r < 20, \quad (2)$$

$$(u - g - 0.82)^2 + (g - r - 0.30)^2 + (r - i - 0.09)^2 +$$

$$(i - z - 0.02)^2 < 0.04,$$

$$\text{Type}_{SDSS} = 6 \text{ (star)}.$$

This selection volume is essentially a 0.2 magnitude four color sphere around the position of F stars in color space (Ivezić et al. 2007). F stars are useful standards because they are common, and because their luminosity peaks roughly in the middle of our $g_{P1}r_{P1}i_{P1}z_{P1}$ wavelength range.

We examine the reduced χ^2 distribution for single filter F star light curves, assuming a constant luminosity model, i.e.:

$$\chi^2_{\text{red}} = \frac{1}{n-1} \sum \frac{(m_i - \bar{m})^2}{\sigma_i^2} \quad (3)$$

$$\bar{m} = \frac{\sum m_i / \sigma_i^2}{\sum 1 / \sigma_i^2}.$$

The quantity χ^2_{red} should approach unity for large ensembles of constant sources, implying that the variation in the mean magnitude is consistent with the error bars. We plot the median χ^2_{red} versus the average of the error bars from different

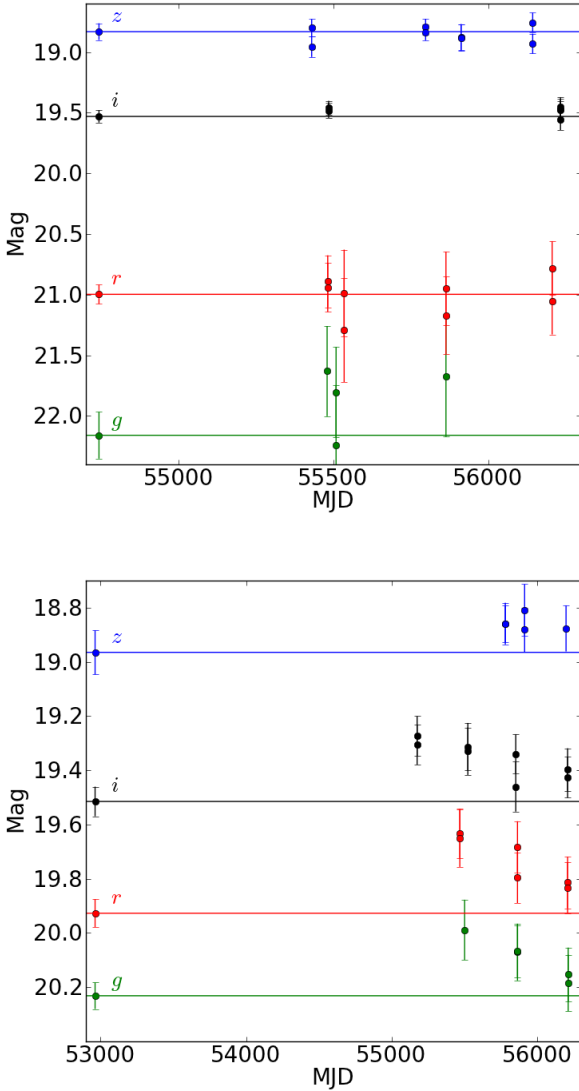


FIG. 2.— Typical SDSS-PS1 light curves from a Stripe 82 photometric standards (top) and variable objects (bottom). The four lines in each figure represent, from bottom to top, the light curve from g (in green), r (in red), i (in black) and z (in blue) filters. The first data point on each light curve and the horizontal line are taken from (PS1-converted) SDSS, and all other datapoints are from PS1.

measurements in Fig. 3. χ_{red}^2 is never 1, but it is fairly constant with respect to the size of the error bars (although there is a small positive correlation between the two). The square root of this constant is 1.387, 1.327, 1.249, 1.228 and 1.170 in g_{P1} , r_{P1} , i_{P1} , z_{P1} and y_{P1} filters, respectively. We multiply the standard PS1 error bars by these constants in our work.

4. TDSS VARIABILITY MEASUREMENT

TDSS aims to take full advantage of the SDSS-PS1 combined dataset to select a highly pure sample of variable objects without any overt bias with regard to color or variability pattern. To achieve this goal we preselect targets whose variability can be robustly measured. We then combine data across filters in which a given source is well-measured into a single three dimensional parameter space. Finally, we use a kernel density estimator (KDE, Rosenblatt 1956; Parzen 1962) and

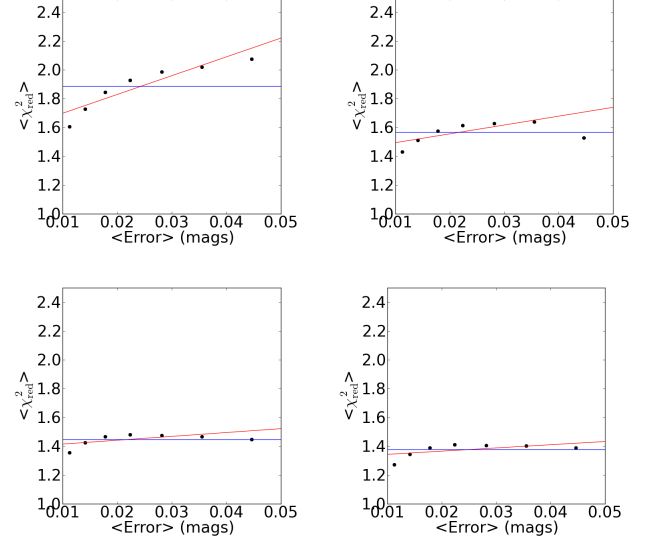


FIG. 3.— χ_{red}^2 versus average photometric measurement error for F stars that satisfy Eq. 2 in g_{P1} (upper left), r_{P1} (upper right), i_{P1} (lower left) and z_{P1} (lower right). Each is fit as a constant (blue) and as a line (red). The data are very roughly consistent with a constant model, with a different constant for each filter.

a Stripe 82 training set to assign each object a probability of being a true variable object based on its location within this 3D KDE space. Fig. 4 outlines this process.

Among objects detected in both SDSS and PS1, we preselect a subset with good data quality to avoid wasting computational resources on sources for which we could not reliably measure variability at the 0.1 magnitude level by requiring that

$$\begin{aligned}
 &17 < i < 21, \\
 &g, r, z > 16, \\
 &\text{types}_{\text{SDSS}} = 6 \text{ (star)}, \\
 &r_{22} > 5'', \\
 &r_{17} > 10'', \\
 &r_{15} > 20'', \\
 &r_{13} > 30'', \\
 &n_{\text{PS1 } griz} > 10.
 \end{aligned} \tag{4}$$

Here, all magnitudes are SDSS PSF magnitudes. We find that 95% of such objects have SDSS i band errors and PS1 i_{P1} mean errors of less than 0.1 magnitudes at $i = 21$. The $i > 17$ and $g, r, z > 16$ requirements prevent selection of very bright sources whose flux would bleed into neighboring spectroscopic fibers. We are obtaining spectra of $16 < i < 17$ targets with smaller telescopes and will discuss this bright extension of TDSS in a future paper. We restrict ourselves to unresolved objects (SDSS morphological type "star"), because it is difficult to perform consistent measurements of extended sources in varying observation conditions at the precision we need. Based on our experience with visual inspection, we also require that the sources not have an $i < 22$ neighbor within $5''$ as this can confuse the photometry ($r_{22} > 5''$). Similarly, we require no $i < 17$, 15, 13 neighbors within $10''$, $20''$, $30''$, respectively. Finally, we require PS1 detections at more than 10 epochs across the $g_{\text{P1}}r_{\text{P1}}i_{\text{P1}}z_{\text{P1}}$ filters ($n_{\text{PS1 } griz} > 10$) for each object to ensure that we have a significant amount of variability information. This last requirement is the most restric-

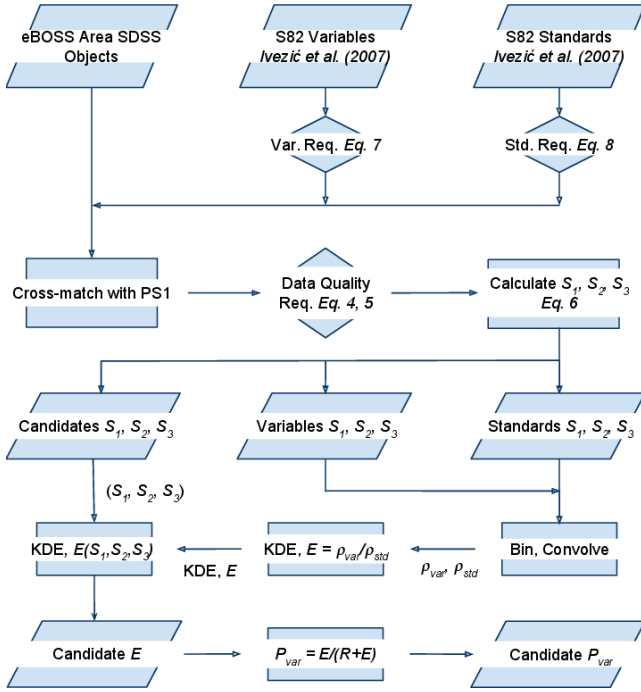


FIG. 4.— A flowchart describing the process by which we determine P_{variable} , the probability that a given candidate is a variable object. Parallelograms represent data objects. Diamonds are conditional statements that reject some of the data. Rectangles are functions. We start with all SDSS objects in the eBOSS area as well as two Stripe 82 training sets, Variables (both stellar and AGN) and Standards. Each of the three data sets pass a set of data quality cuts and are cross-matched with PS1 data. Using PS1 data and single epoch SDSS data, we calculate our variability parameters, (S_1, S_2, S_3) for all three datasets. We use the variability parameters from the training sets to produce a Kernel Density Estimate, $E(S_1, S_2, S_3)$. Using this function, which is static across the sky, we can assign every potential TDSS candidate an E value, which is easily converted to P_{variable} , via R , the (assumed constant) ratio of variable objects to nonvariables.

tive and is met by approximately 85% of $17 < i < 21$ SDSS sources with PS1 matches.

We also examine each source in every filter to determine in which filters we can reliably measure variability. For a given source, we only measure variability in filters in which

$$\begin{aligned} \text{err}_{\text{SDSS}} &< 0.1, \\ \text{err}_{\text{PS1}} &< 0.1, \\ n_{\text{PS1}} &> 1. \end{aligned} \quad (5)$$

Here, err_{SDSS} and err_{PS1} are the SDSS and PS1 mean magnitude errors, respectively. n_{PS1} is the number of detections in a single PS1 filter. Because PS1 lacks a u filter and SDSS lacks a y filter, we only examine variability across the $griz$ filters. To eliminate some obvious artifacts, we ignore filters in which the PS1-SDSS difference is greater than 3 magnitudes, unless the PS1-SDSS difference is greater than 1.2 magnitudes in another filter. For a given source, we designate filters which pass these criteria as “good” and only examine the variability of sources with at least two good filters.

Many groups have demonstrated that advanced machine learning algorithms are highly effective at selecting variable objects (Woźniak et al. 2004a; Richards et al. 2011). These algorithms are generally optimized to accept a large number of inputs to distinguish between variable objects and nonvariables with fairly complex routines that can be difficult to

assess. We opted for a simpler 3D KDE estimator for two main reasons. First, TDSS aims to be a variability-only survey, and it is difficult to ensure that machine learning algorithms (which work best with many input parameters) are primarily using variability to select astrophysical variables. For instance, a boosted decision tree, given the $griz$ magnitudes of a set of variable objects (including many quasars) and non-variables (with mainly stars), can locate quasars clustered in color space, which may reproduce quasar color selection and ignore actual variability entirely. Second, the depth and number of observations per source varies significantly across the PS1 survey, and it is difficult to ensure that a complicated machine learning algorithm is operating efficiently across the whole sky when its inputs change. Furthermore, when we restricted a boosted decision tree to a small number of robust parameters that contained no color information, we found that the KDE detected more variable objects at a similar threshold. We discuss our boosted decision tree results in detail in the appendix.

Through extensive testing, we have settled on a simple 3D (S_1, S_2, S_3) KDE parameter space:

$$\begin{aligned} S_1 &= \text{median}(|\text{mag}_{\text{PS1}} - \text{mag}_{\text{SDSS}}|), \\ \text{Var}_{\text{PS1}} &= \text{Variance}_{\text{PS1}} - \text{Err}_{\text{PS1}}^2 (n_{\text{PS1}} - 1), \\ S_2 &= \text{median}(\text{sign}(\text{Var}_{\text{PS1}}) |\text{Var}_{\text{PS1}}|^{1/2}), \\ S_3 &= \text{median}(\text{mag}_{\text{PS1}}). \end{aligned} \quad (6)$$

Qualitatively, S_1 is the PS1 SDSS difference and represents long term (multi-year) variability. S_2 is the PS1 only variability and represents short term (days to a few years) variability. S_3 is just an apparent magnitude. The word “median” refers to the median magnitude value across all good filters (Eq. 5) for a given source. If there are only two good filters, S_1 and S_2 become minima to prevent individual outlier filters from creating false positive variable targets. All magnitudes used are Point Spread Function (PSF) magnitudes. The PS1 magnitudes, mag_{PS1} , are median magnitudes, used to improve robustness due to the non-simultaneous nature of the PS1 measurements in different filters. Var_{PS1} is an estimate of true PS1 magnitude variability above the expected random variability given the error bars. Var is negative for sources whose photometry randomly varies less than their error bars would indicate. The variable S_2 is a simple function of Var that accounts for this possible negativity while also converting Var into units of magnitudes (where the distribution is more useful for KDE analysis). The variable S_3 is the median PS1 magnitude across good filters. While there is no obvious trend of variability with magnitude, our ability to accurately measure variability decreases as objects get fainter, and using S_3 in our selection allows us to adjust our threshold accordingly.

To assess which bins are the most likely to contain true variable objects, we use a set of confirmed Stripe 82 variable and standard (non-variable) objects. Both catalogs are from Ivezic et al. (2007) and are made with Stripe 82 light curves. Often, our error bars are of order 0.1 magnitudes, so to maintain high purity of the Stripe 82 variable object catalog, we require

$$\begin{aligned} g\text{Ampl} &> 0.1, \\ r\text{Ampl} &> 0.1, \\ i\text{Ampl} &> 0.05, \end{aligned} \quad (7)$$

where Ampl is the estimated amplitude of variation in magnitudes from SDSS. We use a lower threshold in the i band, because both stellar variables and quasars tend to vary less in

redder bands, and because SDSS is shallower and less sensitive to variability in the i band. A total of 89% of Stripe 82 variable objects from Ivezić et al. (2007) satisfy this requirement. To increase the purity of the Stripe 82 standards catalog, we require

$$\begin{aligned} n_{\text{SDSS}} &> 7, \\ \chi_{g \text{ red}}^2 &< 2, \\ \chi_{r \text{ red}}^2 &< 2, \\ \chi_{i \text{ red}}^2 &< 2, \end{aligned} \quad (8)$$

where n_{SDSS} is the number of SDSS measurements and the χ_{red}^2 values are fits assuming a constant magnitude in each filter. Approximately 66% of Stripe 82 standards from Ivezić et al. (2007) satisfy this requirement. To obtain a sample that is similar to TDSS, we use the region of Stripe 82 where $RA > 315^\circ$ or $RA < 60^\circ$. This avoids the $300^\circ < RA < 315^\circ$ region that is at low Galactic latitude and has a stellar density well above that typical in TDSS. Our variable object and standard catalogs have 12,523 and 411,219 sources, respectively.

We divide our standard and variable object KDE spaces into 200 equally spaced bins in each of the three dimensions (i.e., 200^3 total bins). We set the bounds along each dimension to include the middle 99.8% of our variable object training set. The remaining sources are placed in either the minimum or maximum bin as appropriate. We convolve our binned parameter space with a normalized, symmetric Gaussian filter with $\sigma = 5$ bins ($0.02 \times 0.008 \times 0.1$ mags in (S_1, S_2, S_3) -space) so that regions with a small number of sources are filled uniformly as a continuous function. We normalize each KDE density so that it is effectively a probability density.

To prioritize targets, we examine the smoothed, continuous, normalized (so that it integrates to unity) KDE density of Stripe 82 variable objects and standards, which we designate ρ_{var} and ρ_{stan} . We assign each bin in (S_1, S_2, S_3) -space a KDE value, $E(S_1, S_2, S_3)$, defined as

$$E = \frac{\rho_{\text{var}}}{\rho_{\text{stan}}}. \quad (9)$$

Areas of parameter space with the highest values of E are the most efficient places to find variable objects and are initially assigned the highest priority. In Section 5 we will discuss how our final target list does not strictly follow the E value above. This quantity is, in principle, simple to relate to the probability of an object being a variable object:

$$P_{\text{variable}} = \frac{E}{R + E}, \quad (10)$$

Where R is the ratio of nonvariables to variable objects. In practice, R depends on Galactic latitude and longitude, survey depth, observation cadence and the chosen threshold for variability. Different variability surveys could thus have wildly different values of R . Consulting color-based quasar selection, we estimate that an average of 2.4% of sources which pass our data quality preselection in Eqs. 4 and 5 are quasars, which we generally assume to be variable objects. 58% of objects in our Stripe 82 variable object catalog are quasars. We combine these numbers to estimate that approximately 4% of objects which pass our preselection are variable objects. This leads to an estimate $R = 25$, which we use in every region of the sky. While inaccuracies in R will moderately affect our estimates of purity, they do not directly affect the actual targets we select.

5. PRIORITIZATION OF TDSS VARIABLE OBJECTS

Given our allotted fiber density across the sky, we seek a statistically uniformly selected target list of 10 TDSS-only targets deg^{-2} across the entire TDSS area. To move from our 3D "efficiency" space defined in Stripe 82 to this uniform density target list, we divide the sky into equal area "pixels", determine a sensible threshold for our value E (defined in Section 4), accept all targets that cross that threshold in the 20% lowest target density pixels and randomly subsample targets which cross that threshold in the 80% higher target density pixels. Our final sample is then uniform in the sense that objects everywhere pass the same E threshold, but we use more subsampling in denser, low Galactic latitude areas.

We start by dividing the sky into 2×2 degree square pixels. In each pixel, we assign an E threshold that selects exactly 10 TDSS-only targets deg^{-2} after removing the numerous targets shared with the eBOSS CORE quasar program, targets with previous SDSS spectroscopy and a small set of targets selected from the Palomar Transient Factory. We use "TDSS-only targets" to refer to objects selected for observation exclusively by TDSS and refer to the complete set of objects which satisfy our selection criteria as "total targets". We do not formally exclude the objects we share with the eBOSS CORE quasar sample, and they are part of the final TDSS survey. The distinction between these samples is made in our targeting procedure, because TDSS targets that are also in the eBOSS CORE quasar sample are not charged to our survey fiber allotment of 10 targets deg^{-2} .

In Fig. 5 we show three cross sections of our 3D KDE taken from a large region ($135^\circ < RA < 150^\circ$, $45^\circ < DEC < 60^\circ$) for statistical robustness. These cross sections demonstrate how selection varies in S_1 (IPS1-SDSSI) and S_2 (PS1 Variability) at different values of S_3 (median magnitude) where S_1 , S_2 and S_3 are defined in Eq. 6. The three density contours represent the cutoffs we use to obtain 10, 20 and 40 TDSS-only targets deg^{-2} . Our threshold in S_1 and S_2 expands outward at fainter magnitudes indicating, sensibly, that we require stronger variability to observe fainter objects, since they have larger error bars. Objects are generally required to vary by approximately 0.2 magnitudes to meet a 10 target deg^{-2} limitation across most of the sky. Our KDE can fail in regions near the edge of our KDE parameter space where the density of both variable objects and standards is small. To avoid this problem, we assign any object with $S_1 > 0.5$ or $S_2 > 0.25$ a value of $E = 100$ if its E does not already exceed 100. Only 15% of our TDSS-only targets and 8% of our total targets have E assigned to 100.

The KDE that underlies Fig. 5 is derived exclusively from a fixed set of Stripe 82 standards and variable objects and can thus be applied to any area of the sky. However, the positions of the contours in Fig. 5 corresponding to a particular target density are only applicable to a specific 135 deg^2 area of the sky. Different pixels across the sky will have different 10 targets deg^{-2} thresholds (contour positions) corresponding to the variation in density of stellar variables (and stars more generally) across the sky. Fig 6 shows the distribution of 10 target deg^{-2} E thresholds across our pixels. A total of 80% of pixels have a 10 targets deg^{-2} E threshold greater than 45.4, and we adopt this value as our nominal global E threshold. Again, $E(S_1, S_2, S_3)$ is a static function defined by Stripe 82, so the only thing that changes across the sky is the density of objects with $E > 45.4$. Eq. 10 states that the expected variable object

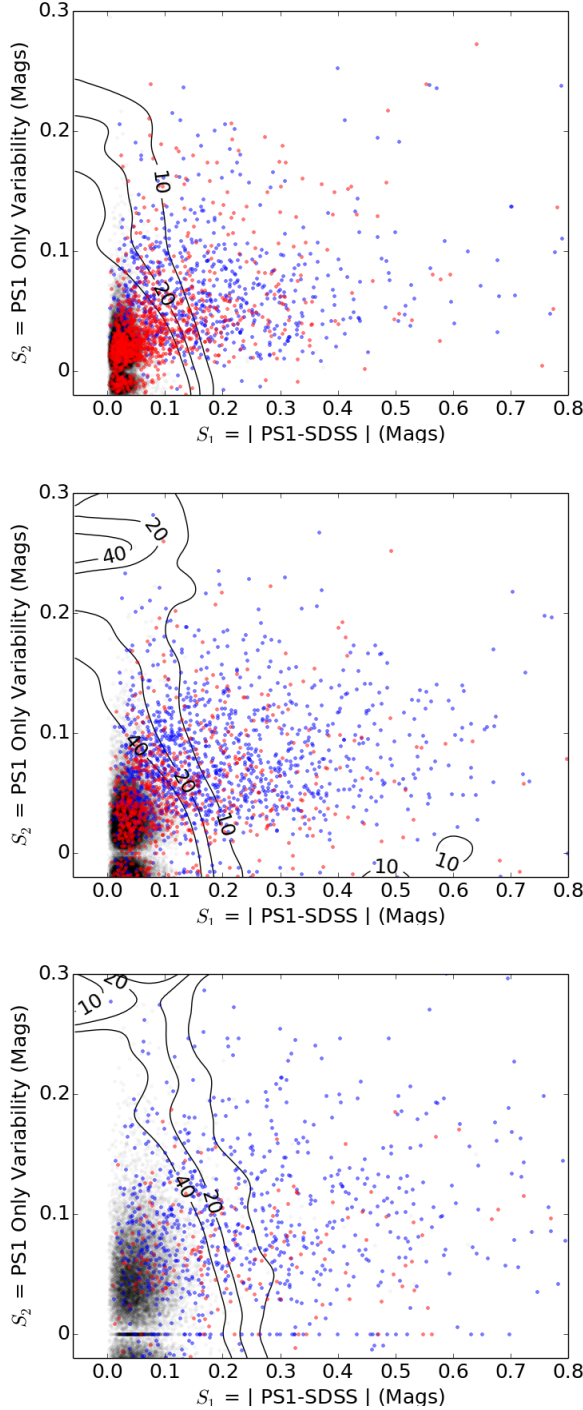


FIG. 5.— 2D cross section from our 3D KDE. The top, middle and bottom panels are cross sections centered around median magnitude, $S_3 = 17.5, 18.5, 20$, respectively. The contour labels specify the number of TDSS-only targets deg^{-2} obtained with each cut. We show the Stripe 82 standards (black), variable quasars (blue) and variable non-quasars (red). Here quasars either have SDSS spectral type ‘q’ or $P(\text{qso}) > 0.5$ according to the eBOSS CORE photometric quasar selection algorithm. The lack of points near the $S_2 = 0$ axis is due to the square root in the definition of S_2 and does not significantly affect the binned KDE.

purity of targets with $E = 45.4$ is 65%. This is a lower bound on our sample purity, and our estimated purity (Section 6) is significantly higher.

Fig. 7 presents the distribution of the density of TDSS-

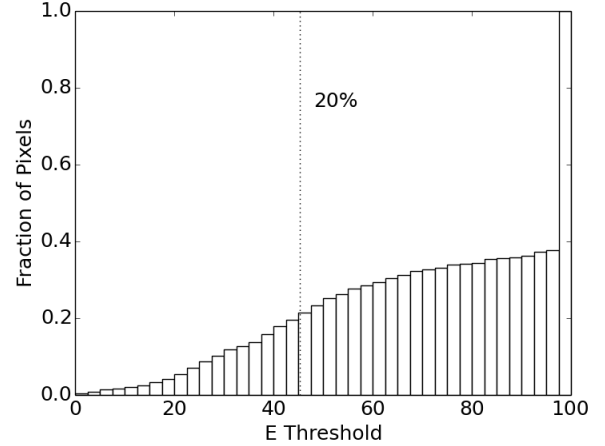


FIG. 6.— The fraction of pixels with a 10 TDSS-only targets deg^{-2} E threshold less than a given value. Our global E threshold, 45.4, is marked with a dotted line. Only 20% of pixels have a 10 targets deg^{-2} threshold less than 45.4. As noted in the text, 15% of TDSS-only targets have their E manually set to 100 which leads to the jump at $E = 100$.

only objects (those objects not selected as part of the eBOSS CORE quasar sample and not having previous spectroscopy) and of all objects that cross the $E = 45.4$ threshold in each pixel. The density rises precipitously in the low Galactic latitude regions at the edges of our survey. This result simply implies that a significant fraction of our variable objects are stars that become more common at lower Galactic latitude. Some of the pixels with low target density are near the very edge of our estimated survey bounds. Many of these pixels will not be included in the actual spectroscopic survey. Globally, the average density of TDSS-only targets with $E > 45.4$ is 14 deg^{-2} , and we are sparse sampling 70% of these sources. The majority of pixels with fewer than 10 targets deg^{-2} with $E > 45.4$, have at least 8 targets deg^{-2} , so only a small number of spectra of $E < 45.4$ objects will be taken.

Having established a threshold, we must still determine how to make a uniform target list with 40 TDSS-only targets (10 deg^{-2}) in each 4 deg^2 pixel. In the 20% of pixels with 40 or fewer TDSS-only targets that cross the E threshold, we simply select the 40 targets with the highest E estimate (a small fraction of which have $E < 45.4$). In the 80% of pixels with more than 40 targets, we prioritize a small number of hypervariable targets (described in Subsection 5.1) and then assign a random priority to the remaining targets with $E > 45.4$, choosing the targets with the highest priority until we reach our 10 deg^{-2} target quota.

Our final step to produce a target list is to visually inspect every object’s SDSS image. Visual inspection was performed by authors Morganson, Green, Anderson and Ruan. Objects judged to have significant flux from nearby neighbors, objects with unflagged processing errors or objects within approximately $30''$ of a diffraction spike are removed. Lower priority objects rise in the queue naturally, with $E < 45.4$ objects being prioritized directly by E value. The fraction of objects removed by visual inspection ranges between 5% and (rarely) 30%. The rejection fraction is highest at low Galactic latitudes where there are many very bright stars (that can influence photometry over distances of several arcminutes) and close stellar pairs (unresolved in the SDSS catalog). Fortunately, these regions also have an abundance of high E targets.

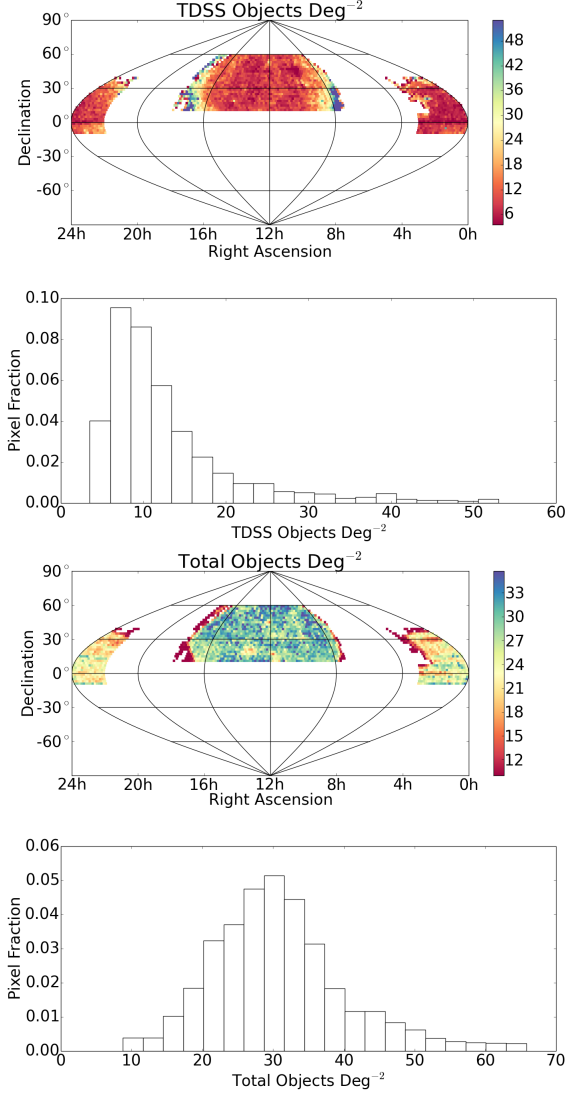


FIG. 7.— Map showing the density of TDSS-only objects (excluding CORE quasars and objects with previous spectroscopy) that exceed the $E = 45.4$ threshold in each 2×2 degree pixel and the distribution of these densities (top 2 panels). Map showing the total density of all objects that exceed the $E = 45.4$ threshold in each pixel, including objects shared with the eBOSS CORE quasar sample and objects with previous spectra, and the corresponding distribution of these densities (bottom two panels).

5.1. TDSS Prioritization of Hypervariables

While the main goal of TDSS is to provide a statistically uniform sample of variable objects, TDSS also provides a unique opportunity to obtain a statistical sample of the most variable objects in the sky, which we designate as hypervariables. While most of these hypervariables would be observed naturally as part of the survey, we wish to ensure that hypervariables which vary above a particular threshold are all observed, regardless of the local target density. Our KDE method is not well-designed to select hypervariables, since the extreme regions of variability space are poorly populated by either variable objects or standards. Instead we reduce our variability parameters from Eq. 6 to a single parameter:

$$V = (\text{median}(|\text{mag}_{\text{PS1}} - \text{mag}_{\text{SDSS}}|)^2 + 4 \text{median}(\text{Var}_{\text{PS1}})^2)^{1/2} \Gamma_1 \\ = (S_1^2 + 4S_2^2)^{1/2}.$$

V_{Thresh}	H	H_{QSO}	H_*	H_{CORE}	H_{prev}	$H \text{ deg}^{-2}$	$H_{\text{low}} \text{ deg}^{-2}$
1.0	8411	7	6489	1274	640	1.05	1.54
1.2	4784	3	4046	481	253	0.60	0.91
1.4	3069	1	2725	224	118	0.38	0.55
1.6	2071	0	1899	120	51	0.26	0.42
1.8	1492	0	1375	80	36	0.19	0.24
2.0	1108	0	1033	52	22	0.14	0.19
2.2	823	0	778	29	15	0.10	0.15
2.4	629	0	595	21	12	0.08	0.10
2.6	483	0	463	12	8	0.06	0.08
2.8	401	0	385	10	6	0.05	0.06
3.0	338	0	323	9	6	0.04	0.06

TABLE 3

Estimated total number of hypervariables (H) that pass different thresholds of V (V_{Thresh} from Eq. 11). We also show the expected number of TDSS-only quasars (H_{QSO}), TDSS-only non-quasars (H_*), quasars shared with the CORE quasar group (H_{CORE}) and targets with previous SDSS spectra (H_{prev}). The last two columns are the total density (deg^{-2}) of hypervariables across the whole survey and the density in a low Galactic latitude region ($120^\circ < \text{RA} < 130^\circ$, $10^\circ < \text{DEC} < 20^\circ$).

This is an elliptical contour of approximately constant density in our (S_1, S_2) variability space. The factor of 4 accounts for the fact that x , the SDSS-PS1 difference, is generally of order twice y , the PS1 only variability. This ratio is not exact and is specific to this data. It is likely due to the longer time scales of the SDSS-PS1 difference.

Table 3 lists the number of hypervariables as a function of different thresholds of V (Eq. 11). Using the densities presented here and the density map shown in Fig. 8, we set our V threshold to 2.0 magnitudes. This choice yields 1,108 targets, most of which would have likely been observed naturally by our KDE selection method. Globally, this population density is 0.14 deg^{-2} but in a representative low Galactic latitude region ($120^\circ < \text{RA} < 130^\circ$, $10^\circ < \text{DEC} < 20^\circ$), the density of hypervariables is 0.19 deg^{-2} , 2% of our low latitude targets. While the majority of targets TDSS selects have colors consistent with being quasars, approximately 95% of our hypervariables do not. We will briefly investigate the likely identities of hypervariables in Section 8.3.

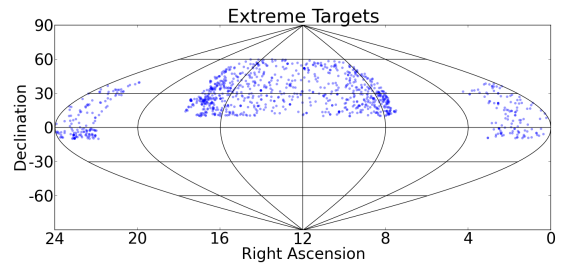


FIG. 8.— Map showing the locations of $V > 2.0$ hypervariables across the sky in equatorial coordinates.

6. ANTICIPATED PURITY

When evaluating our selection criteria, we were primarily concerned with the purity of all targets selected at a given threshold, P_{tot} , and the purity of our TDSS-only target list after eBOSS CORE quasars and targets with previous SDSS spectroscopy are removed, P_{tar} . We estimate the anticipated purity of our total sample using our results in Stripe 82.

Specifically, building on Eq. 10, we define:

$$P_{\text{tot}} = \frac{f_{\text{var S82}}}{R f_{\text{stan S82}} + f_{\text{var S82}}}, \quad (12)$$

$$R = 25,$$

$$P_{\text{tar}} = \frac{f_{\text{tar}} f_{\text{var S82}}}{R f_{\text{stan S82}} + f_{\text{tar}} f_{\text{var S82}}}, \quad (13)$$

$$f_{\text{tar}} = \frac{N_{\text{tar}}}{N_{\text{tar}} + N_{\text{CORE}} + N_{\text{prev}}}.$$

Here, $f_{\text{var S82}}$ and $f_{\text{stan S82}}$ are the fraction of Stripe 82 variable objects and standards (as defined in section 4) that pass a given threshold. The quantity R is the expected ratio of nonvariables to variable objects discussed in Section 4, and f_{tar} is the fraction of objects which pass a given threshold that will be TDSS-only targets. N_{tar} , N_{CORE} and N_{prev} are the numbers of TDSS-only targets, eBOSS CORE quasar targets and objects with previous SDSS spectroscopy that pass a given threshold, respectively. The sum $N_{\text{tar}} + N_{\text{CORE}} + N_{\text{prev}}$ is the total number of objects that pass a given threshold.

To assess the performance of our variable object selection algorithm, we tested it on a large, representative patch of sky, the $135^\circ < RA < 150^\circ$, $45^\circ < DEC < 60^\circ$ region previously mentioned in Section 4. We used this region to set our selection threshold to obtain 10, 20 ... 60 TDSS-only targets deg^{-2} in Table 4. Having set a threshold to obtain a known density of targets, cross-matched with eBOSS and SDSS databases to remove eBOSS CORE quasars and objects with previous spectra and calculated purity with Eq. 13, we can estimate the number of low variability sources that scatter into our selection space:

$$N_{\text{lovar}} = N_{\text{tar}}(1 - P_{\text{tar}}). \quad (14)$$

To estimate our quasar fraction, we assume that everything in the color box:

$$\begin{aligned} u_{\text{SDSS}} - g_{\text{SDSS}} &< 0.8, \\ g_{\text{SDSS}} - r_{\text{SDSS}} &< 0.65 \end{aligned} \quad (15)$$

is a quasar.

There are

$$N_* = N_{\text{tar}} - N_{\text{QSO}} - N_{\text{lovar}} \quad (16)$$

remaining objects which are expect to be mostly stellar variable objects. We also calculate N_{CORE} , N_{prev} and N_{tot} .

Table 4 demonstrates how various quantities change as we lower our target selection threshold. The thresholds are set so that the 20th percentile pixel (as described in Section 5) has $N_{\text{tar}} 20$. We actually derive our statistics from our larger test region, which is quite similar to the 20th percentile pixel (the density of targets in this region is $N_{\text{tar test}}$). We acquire 10 TDSS-only spectra deg^{-2} , so the final line is most useful. The TDSS selection algorithm at that surface density produces a target list that is $P_{\text{tot}} = 95\%$ pure. Many of these targets are shared with the eBOSS CORE quasar sample or have previous SDSS spectra. After these targets are removed, the TDSS-only targets are $P_{\text{tar}} = 88\%$ pure. Note that "low variability" sources are sources that did not vary in the Stripe 82 data. Some unknown, but likely significant, fraction of these sources are true variable objects that varied during the PS1 epochs or between SDSS and PS1. In addition, our visual inspection removes a significant fraction of non-variable objects that is not accounted for in this analysis.

Since most of the quasars we select are shared with the eBOSS CORE quasar sample, the TDSS-only targets are approximately 90% non-quasars (mostly variable stars). In practice, we expect to find a significant fraction of unusual quasars with colors not described by Eq. 15, so precise estimates of the quasar fraction will require spectra. Our data and selection method are optimized for 10 TDSS targets deg^{-2} . If we were to expand our target list to the 20 targets deg^{-2} threshold, we would be selecting 6.5 additional variable objects and 5.6 additional standards in our test field (roughly 5.4 variable objects and 4.6 standards in our 20% field). Our additional targets would be only 54% pure. Selecting a "deeper" set of variable objects with high purity likely requires higher precision PS1 data or significantly better-sampled light curves.

7. TDSS SELECTION OF *I*-BAND DROPOUTS

TDSS strives to produce a sample of variable objects that is unbiased in color space. We make one small exception for *i*-dropouts, objects that are observed in the *z* band but are either not observed in any bluer bands or have extremely large *i* - *z* colors. Among known astrophysical *i*-dropouts are late M and L-type dwarfs and $z \approx 6$ quasars, all of which are rarely detected and may have interesting variability properties. Our two filter requirement would exclude these objects if we did not create a separate pipeline to identify them.

Our *i*-dropout selection method closely follows the main selection method described in Section 4. First, we make an initial database level cut:

$$\begin{aligned} i_{\text{SDSS}} - z_{\text{SDSS}} &> 1.0 \\ \text{err}_{z_{\text{SDSS}}} &< 0.1, \\ \text{err}_{i_{\text{SDSS}}}, \text{err}_{r_{\text{SDSS}}}, \text{err}_{g_{\text{SDSS}}} &> 0.1, \\ r_{22} &> 5'', \\ r_{17} &> 10'', \\ r_{15} &> 20'', \\ r_{13} &> 30'', \\ n_{\text{PS1 } z}, n_{\text{PS1 } y} &> 3. \end{aligned} \quad (17)$$

The first three requirements are all purely SDSS-based and are designed to find *i*-dropouts while excluding any sources found in the main sample. The next four requirements remove objects whose photometry has likely been altered by a nearby bright object. The final requirement ensures that we have sufficient PS1 data to make a variability measurement. These criteria yield 11,594 sources with typical limiting magnitudes of $z < 19.9$, $y < 19.7$. These requirements also ensure that the objects are real and not just cosmic rays or other artifacts in a single *z* band image, which can be problematic for *i*-dropout searches (Fan et al. 2001; Morganson et al. 2012).

To select long term variable objects, we would naturally wish to use the SDSS-PS1 *z* magnitude difference. Our filter transformations in Eq. 1, however, are not designed to work with *i*-dropouts, which are bound to have extreme colors, so we must derive our own SDSS-PS1 filter corrections. In Fig. 9, we fit $\Delta z = z_{\text{SDSS}} - z_{\text{PS1}}$ versus $zy = z_{\text{PS1}} - y_{\text{PS1}}$ as a line, $\Delta z = a + b zy$, by minimizing the absolute deviations:

$$S = \sum_i \left| \frac{\Delta z_i - (a + b zy_i)}{\sigma_i} \right|, \quad (18)$$

$N_{\text{tar } 20}$	$N_{\text{tar test}}$	N_{QSO}	N_*	N_{lovar}	N_{CORE}	N_{prev}	N_{tot}	P_{tar}	P_{tot}
60	67.8	3.3	16.7	47.7	18.6	27.4	113.7	45.2	58.0
50	56.5	3.0	16.6	36.8	18.0	25.8	100.3	49.2	63.3
40	45.4	2.7	16.4	26.4	17.2	24.2	86.8	54.5	69.6
30	35.2	2.4	15.8	17.0	16.2	22.4	73.8	61.4	76.9
20	23.7	2.1	14.1	7.5	14.6	19.8	58.1	73.3	87.1
10	11.3	1.4	8.3	1.7	10.8	14.6	36.7	86.4	95.4

TABLE 4

Estimated target counts and purities from Stripe 82 tests at different variability cutoffs. All counts are deg^{-2} . All purities are percentages. $N_{\text{tar } 20}$ is the number of targets in the 20th percentile pixel for a given threshold while $N_{\text{tar test}}$ is the number of targets in our test field. N_{QSO} , N_* and N_{lovar} are the estimated numbers of TDSS-unique quasars, stars and low-variability objects, respectively. N_{CORE} and N_{prev} are the estimated numbers of objects we share with the CORE quasar sample or have previous SDSS spectroscopy. N_{tot} is the total number of candidates. P_{tar} and P_{tot} are the estimated purities of our TDSS-only targets and our total targets, respectively.

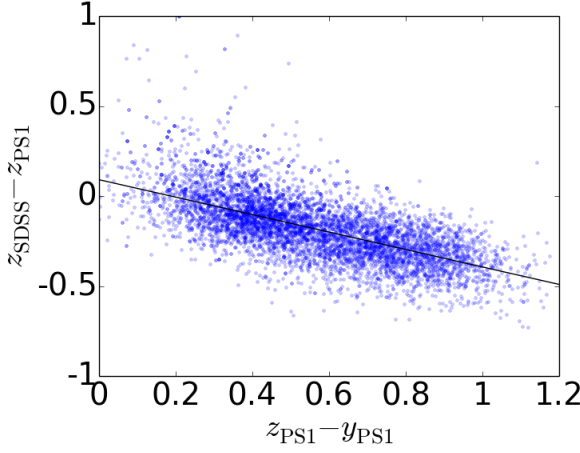


FIG. 9.— The $z_{\text{SDSS}} - z_{\text{PS1}}$ differences of i -dropouts versus their $z_{\text{PS1}} - y_{\text{PS1}}$ colors. The line was defined using a linear minimum absolute deviations fit, and we use it to compare individual z_{PS1} 's to z_{SDSS} 's.

yielding

$$\begin{aligned} a &= 0.141, \\ b &= -0.525. \end{aligned}$$

Minimizing the absolute deviations is more robust to outliers than a typical χ^2 method. With this linear fit, we can define an expected z_{SDSS} given PS1 colors:

$$z_{\text{SDSS}}^* = z_{\text{PS1}} + a + b(z_{\text{PS1}} - y_{\text{PS1}}). \quad (19)$$

$z_{\text{SDSS}}^* - z_{\text{SDSS}}$ is 0 for a typical i dropout in our sample. With this correction, we can define a 2D KDE parameter space analogous to the first two dimensions of our main selection KDE in Eq. 6

$$\begin{aligned} S_1 &= |z_{\text{SDSS}}^* - z_{\text{SDSS}}|, \\ \text{Var}_{\text{PS1}} &= \text{Variance}_{\text{PS1}} - \text{Err}_{\text{PS1}}^2 (n_{\text{PS1}} - 1), \\ \text{Var}_{zy \text{ PS1}} &= 0.5(\text{Var}_z \text{ PS1} + \text{Var}_y \text{ PS1}), \\ S_2 &= \text{sign}(\text{Var}_{zy \text{ PS1}}) |\text{Var}_{zy \text{ PS1}}|^{1/2}. \end{aligned} \quad (20)$$

Figure 10 shows our 2D KDE variable i -dropout selection space. The objects which satisfy the criteria in Eq. 17 tend to be at the faint end of their selection space with magnitude errors near the 0.1 magnitude limit. Their distribution is correspondingly more broad than that of the sources in Figure 10. We lack a large sample of confirmed variable i -dropouts and therefore cannot produce a training set as we did for the main population. Instead, we select the 5% outliers in variability space. There is a small population of sources with negative

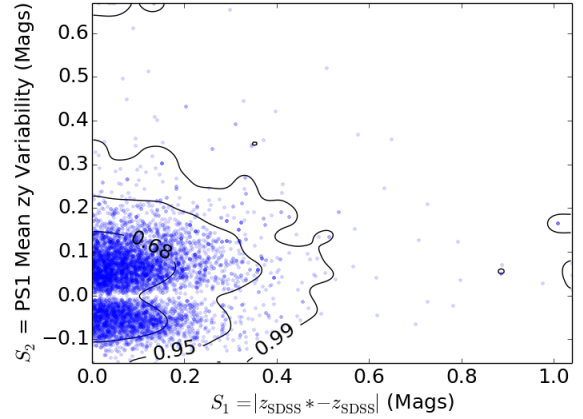


FIG. 10.— The locations of the i -dropouts in the 2D KDE space defined in Eq. 20. We selected variable targets from outside the 95% contour.

PS1 variability (in which the standard deviation is less than what one would expect from the error bars as described in Section 4) and only moderate PS1-SDSS difference. To avoid "rewarding" sources for having negative PS1 variability, an area of parameter space that is rare, but not particularly likely to indicate true variability, we eliminate sources that satisfy

$$S_1 < 0.6, S_2 < 0, \quad (21)$$

where S_1 and S_2 are defined in Eq. 20. In total, 221 i -dropouts satisfy our selection criteria. Of these, only 73 pass our visual inspection and are included in the TDSS target list.

Only seven previously discovered $z \approx 6$ quasars (Fan et al. 2001, 2006; Morganson et al. 2012; Bañados et al. 2014) satisfy our initial selection criteria, and only one passes our variability threshold. This is not entirely surprising as cosmological time dilation will significantly reduce any observed variability from these quasars. In addition, MacLeod et al. (2010), Morganson et al. (2014) and others have found significant anticorrelation between quasar variability and luminosity, and $z \approx 6$ quasars detected by SDSS are necessarily extremely luminous.

8. PHOTOMETRIC CLASSIFICATION OF ALL TDSS TARGETS

The algorithm described in Sections 4 and 5 produces a target list that includes 242,513 objects. This list has approximately 10% more sources than we will be able to target spectroscopically due to a combination of extra area and extra density. Nevertheless, the fractions of different classes of objects in this list should closely resemble the final spec-

troscopic sample. While we do not make any explicit use of color in our variable object selection, classifying our objects by color will allow us to anticipate our final results. We do not attempt to correct for Milky Way reddening in our photometry, because colors do not directly influence our selection. Since the TDSS area is at high Galactic latitude, this only introduces a small error on our color measurements.

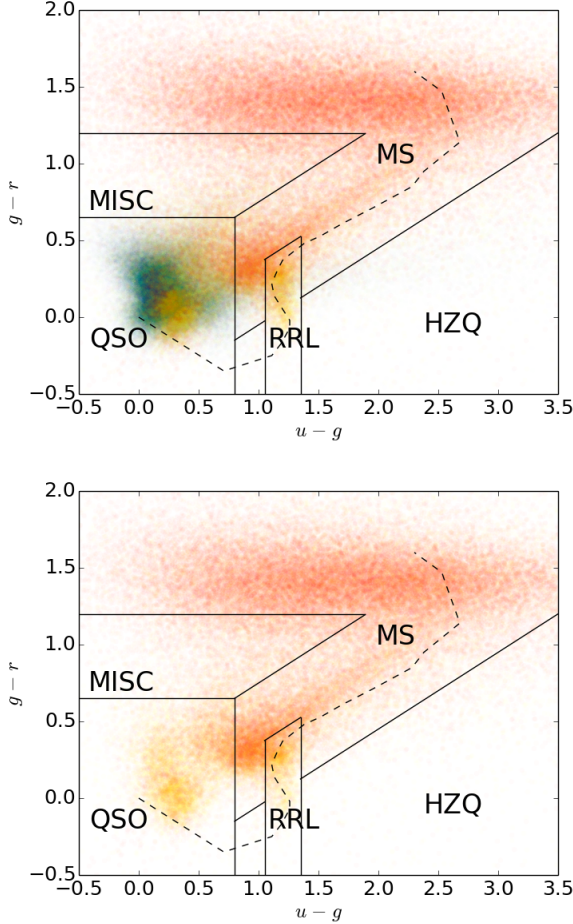


FIG. 11.— The SDSS $g-r$ versus $u-g$ distribution of all TDSS targets (top) and TDSS-only targets after we remove CORE quasars and objects with previous SDSS spectroscopy (bottom). Low (high) priority eBOSS color-selected quasars are in green (blue). Low (high) priority non-quasars are in red (yellow). The QSO, MS, RRL and HZQ regions are the areas of color space that contain most quasars, main sequence stars, RR Lyrae stars and high-redshift ($z > 2.5$) quasars, respectively. The dotted line represents the stellar main sequence. The horizontal blur at $g-r > 1.5$ is due to objects not being detected in the u band and being assigned an essentially random ‘Lupitude’. Similarly, marginal u detections are biased to lower Lupitudes and our $u-g$ distribution is shifted lightly to the left of the fiducial main sequence line.

In Fig. 11, we show the SDSS $g-r$ versus $u-g$ distribution of all TDSS targets. Here, we include all objects with previous spectroscopy as well as those objects we share with the eBOSS CORE quasar group. The extended horizontal cloud at $g-r > 1.2$ is mostly due to objects with essentially zero flux in the u band. These objects have very large error bars in $u-g$. We define regions of color-space on the plot (using

Category	All Targets		TDSS-only Targets	
	N _{objects}	% of Total	N _{objects}	% of Total
MS	75754	31.2	67922	71.5
QSO	143052	59.0	12754	13.4
RRL	7358	3.0	4384	4.6
HZQ	6948	2.9	3059	3.2
MISC	9401	3.9	6889	7.3

TABLE 5
Numbers and fractions of different broad color-based categories as shown in Fig. 11 in our total sample and our TDSS-only sample.

SDSS colors):

$$\text{QSO} : u-g < 0.8, g-r < 0.65, \quad (22)$$

$$\text{RRL} : u-g < 1.35, u-g > 1.05, \\ g-r < 0.5(u-g) - 0.15,$$

$$\text{MS} : g-r > 1.2 \text{ or} \\ u-g > 0.8, g-r < 0.5(u-g) + 0.25, \\ g-r > 0.5(u-g) - 0.55, \text{ not RRL},$$

$$\text{HZQ} : u-g > 0.8, g-r < 0.5(u-g) - 0.55, \text{ not RRL},$$

$$\text{MISC} : g-r > 0.6, g-r < 1.2, g-r > 0.5(u-g) + 0.25.$$

Our categories are named to indicate the primary type of expected variable object in each region, but no category will be absolutely pure. The QSO fiducial color region is mostly quasars and other AGN and is identical to that defined by our quasar criteria in Eq. 15. RRL contains RR Lyrae stars and other variable F stars. MS contains the bulk of the main sequence. HZQ is the region where high-redshift ($z > 2.5$) quasars typically reside. There is no dominant astrophysical identity of the MISC (miscellaneous) sources, but various white dwarf binary systems are included. Consistent with previous variability studies, the region with the most targets is the QSO region (59.0%) followed by the MS (31.2%) as shown in Table 5. To estimate our total number of quasar candidates, we add our QSO and HZQ objects and subtract 10% to obtain 135,000. We take 90% of the sum of our other three categories to estimate 85,000 stellar variables.

We can make more sophisticated color classifications of particular classes of objects. Using eBOSS CORE quasar color-based photometric classification and previous SDSS spectroscopy, we can alternately define “CORE quasars” as those objects for which

$$P(\text{qso}) > 0.5 \text{ or} \quad (23)$$

$$\text{Class}_{\text{SDSS}} = \text{QSO}.$$

$P(\text{qso})$ is provided by the eBOSS CORE quasar team (Myers et al. 2015, *in preparation*) which uses the *XDQSOz* algorithm (Bovy et al. 2012) and $\text{Class}_{\text{SDSS}}$ is the SDSS spectral class from previous spectroscopy. These criteria do not include the small number of potential quasars selected exclusively by the PTF variability quasar search. The eBOSS quasar classifier actually only applies to $z > 0.9$ quasars, but most lower redshift quasars are either swept up into this classifier or already have previous spectra. In Fig. 11, the 134,289 CORE quasars (as now defined by Eq. 23) are shown in green and blue, with the highest priority variable objects in blue. The 108,224 objects not satisfying Eq. 23 are shown in red and yellow with the highest priority objects shown in yellow.

Reassuringly, the bulk of our quasars are centered around $u-g = 0.2$, $g-r = 0.2$, the known center of the quasar locus. There is also a high density of points along the main

sequence, although there is more than the ≈ 0.1 magnitude of scatter we would expect from our statistical error bars. This just indicates that many of our variable stars have somewhat unusual colors and is to be expected for variables (e.g. unresolved binaries or stars with particularly active photospheres). One notable subpopulation of this plot is the "blue cloud" of 7,548 sources not classified as quasars by Eq. 23 in the region defined by

$$\begin{aligned} 0.5 < u-g < 1.0 \\ 0.1 < g-r < 0.5. \end{aligned} \quad (24)$$

This cloud extends off the left of the main sequence and while photometrically blue, is colored red in our plot. A large fraction of these objects are likely to be $z \approx 2.8$ quasars. This is a well-known region in color space where quasars begin to overlap with the main sequence and the color selection used to produce the eBOSS CORE quasar sample is insufficient to distinguish the two. The addition of variability information has likely allowed us to break the color degeneracy and may be used in the future to extend the redshift range of quasar samples. Note that the coloring in Fig. 11 is effectively opaque in high density regions with non-quasars being plotted over quasars. Underneath the "blue cloud" there are also 11,160 objects identified as quasars by the criteria in Eq. 23 "underneath" the "blue cloud" in the top figure.

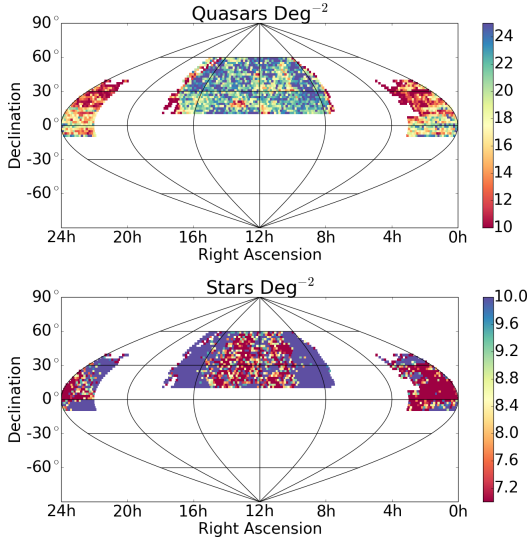


FIG. 12.— The estimated density of quasars deg^{-2} (top) and stars deg^{-2} (bottom) in our target list across the sky. The "ratty" edges are in very dusty regions that are not actually part of our sample.

In Fig. 12 we show separately the estimated density of quasars and stars deg^{-2} across our target sample. In this plot, we define quasars via the simple color box in Eq. 15. Our map does not perfectly match the eBOSS area and some extra areas near the edges contain significant (unaccounted for) dust that limits depth and reddens quasars out of our color box. This reddening, combined with geometric incompleteness near the edges of our survey, lower densities near the edge of our field. Beyond these small underdense edges that will not be included in the final survey, our targets are uniformly distributed, not displaying the strong Galactic density variation of the sky plots in Fig. 7.

Fig. 13 shows the magnitude distribution of the TDSS targets. In general, we would expect unbiased magnitude dis-

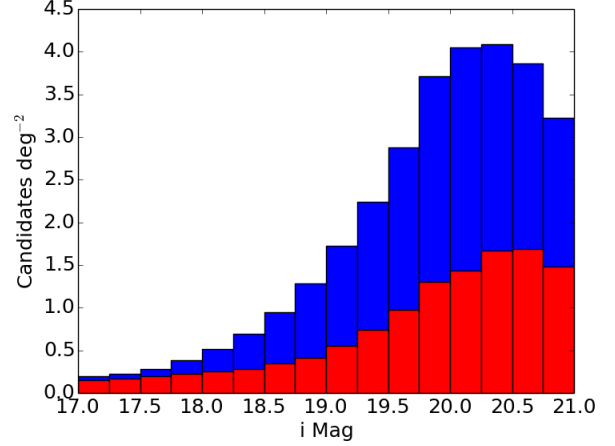


FIG. 13.— The magnitude distribution of all targets (blue) and TDSS-only targets (red).

tributions to increase exponentially at fainter magnitudes. Instead, our magnitude distribution peaks at $i = 20.25$. This is a price we pay to ensure high purity. The requirement of detections in multiple filters, our accounting for error bars in Eq. 5 and the increased variability requirements at fainter magnitudes as shown in Fig. 5 all decrease the target density at fainter magnitudes. Our TDSS-only targets have a larger tail on the bright end than the CORE quasars and objects with previous spectra. This result can be explained by the fact that our variable objects are mostly stars (see Table 4), and stars are more concentrated at brighter magnitudes relative to quasars.

8.1. The Quasar Population

We can probe our likely quasar targets in significantly more detail using a combination of previous spectroscopy and photometry. We are particularly interested in seeing if we are strongly biased towards selecting quasars in a particular color or redshift region. If this were the case, it might indicate that our filter transformations in Eq. 1 were failing catastrophically in that region. Fortunately, as we show below, the only redshift and color biases are subtle and expected.

Fig. 14 shows the redshift distribution of three categories of spectroscopic quasars: all the unresolved, $17.8 < i_{\text{SDSS}} < 19.1$ SDSS spectroscopic quasars in the TDSS footprint, those with $P_{\text{qso}} > 0.5$ according to the eBOSS CORE quasar sample and those that make our target list. We chose these limits because the eBOSS CORE bright limit is 17.8, and the previous SDSS spectroscopic faint limit (for the main $z < 2.5$ quasar population) is approximately 19.1. To be clear, these quasars all have previous SDSS spectroscopy and will not generally be reobserved in TDSS. The eBOSS team excludes $z < 0.9$ quasars from their sample. In general, TDSS recovers 30% of all spectroscopically-confirmed quasars across a broad range of redshift. There are no sharp gaps or spikes that indicate that quasars at particular redshifts are being over-selected or under-selected due to Eq. 1 or other effects.

The bottom panel of Fig. 14 compares the selection efficiency of the CORE quasar sample and TDSS. TDSS underselects $z < 0.2$ objects spectroscopically classified as quasars. Most lower redshift objects with SDSS spectral classification of 'QSO' are in fact lower luminosity active galaxies whose emission is not dominated by the central black hole. This is indicated by the fact that $0.2 < z < 2.5$ quasars from the plot have mean (median) $u-g$ color of 0.22 (0.25), whereas the

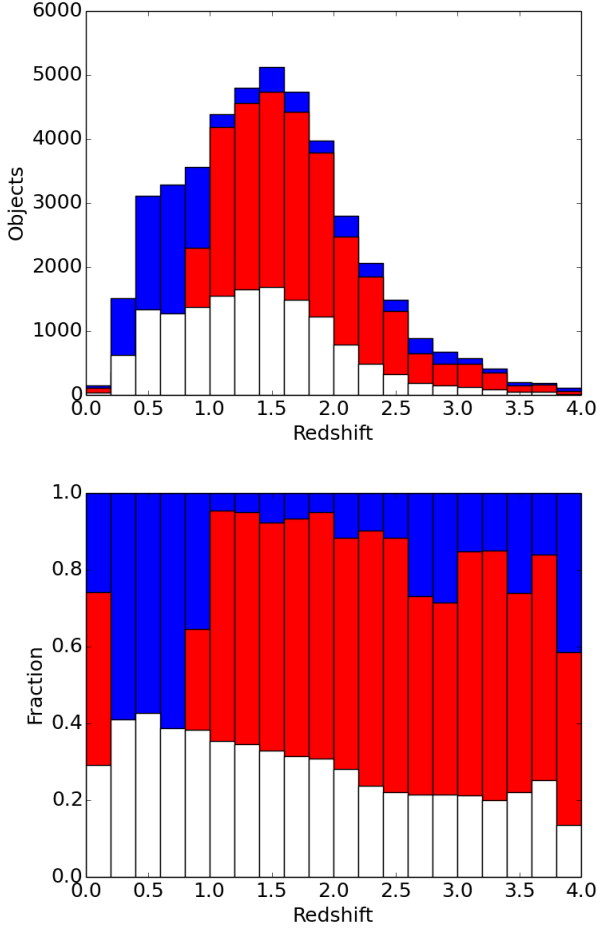


FIG. 14.— The redshift distribution of quasars with previous SDSS spectroscopy (top). The histograms show all unresolved, $17.8 < i_{\text{SDSS}} < 19.1$ spectroscopic quasars in the TDSS area (blue), spectroscopic quasars that have an $XDQSOz$ probability $P_{\text{qso}} > 0.5$ according to the CORE quasar team (red) and quasars that make our final target list (white). The bottom panel shows the fraction of each population as a fraction of the total spectroscopic quasar population. Note that the $XDQSOz$ probability used to select eBOSS CORE quasars intentionally excludes $z < 0.9$ quasars from their sample.

$z < 0.2$ quasars in this plot have mean (median) $u-g$ color of 0.95 (0.53). This extra redness is indicative of significant host galaxy flux contamination. Both the CORE quasar sample and the TDSS sample have a decreasing selection efficiency with increasing redshift. For the CORE quasar sample, this effect arises because quasars have less distinct colors at $z > 2.5$, particularly at $z \approx 2.8$ where quasars have similar optical colors to main sequence stars. The TDSS roll-off in efficiency is more gradual and is likely due to the fact that higher redshift quasars vary more slowly due to cosmological time dilation as well as their high luminosities and implied large black hole masses.

Fig. 15 compares $g-r$ versus $u-g$ for all $17.8 < i < 21.0$, $P_{\text{qso}} > 0.5$ CORE quasars and $17.8 < i < 21.0$ spectroscopically identified quasars as well as the subset of those quasars selected by TDSS. The distributions are qualitatively nearly identical. Fig. 15 (bottom) shows the ratio of the two populations across color space. Across the main quasar locus, TDSS recovers 20-30% of the CORE and spectroscopic quasars. In Fig. 15 (top) there is a faint peninsula of CORE quasar targets stretching from $u-g = 0$, $g-r = -0.2$ to $u-g = -0.5$, $g-r = -0.5$ that are not selected by TDSS in Fig. 15 (middle). These

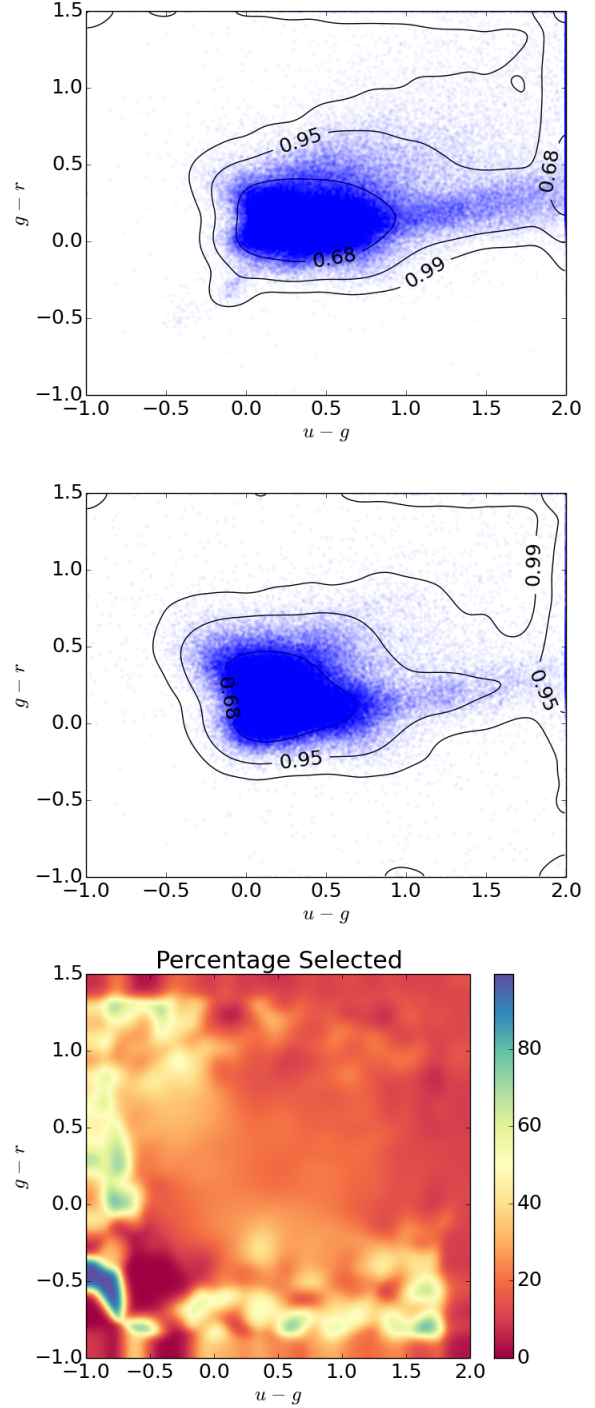


FIG. 15.— The SDSS $g-r$ versus $u-g$ distribution of all quasars for all $17.8 < i < 21.0$ quasars in the eBOSS area (top) and the subset of those selected by TDSS (middle). The bottom panel shows the ratio of the two populations.

objects are likely to be white dwarfs. Excluding this area, TDSS shows a broad tendency to be more complete at the blue end in both the $u-g$ and $g-r$ axes, although there is a low completeness region in the lower left hand corner of Fig. 15 (bottom) that may be due to small number statistics. This preference for blue objects may partly stem from our decreasing completeness at higher redshift shown in Fig. 14. For a given i , we will also generally be more sensitive to variability

for blue objects that are bright in r and g . So our i magnitude limit may lead to an implicit blue source selection bias.

It is not surprising that the CORE quasar team is significantly more complete at selecting quasars than we are. Their selection is focused on quasars, and it is roughly 4 times larger than our sample. But it should be noted in the analysis above, we do not (and cannot) evaluate the fraction of quasars selected by their variability with TDSS that are missed by conventional color selection. Some poorly constrained fraction of quasars are reddened by dust or otherwise have non-standard colors, and the spectra from TDSS will allow us to study how well many of these quasars we can select from their variability.

8.2. The Stellar Population

Using spectroscopy and eBOSS color-based quasar selection, we can statistically remove most quasars from our sample and investigate the colors of our stellar targets. Again, TDSS does not select stellar targets with color classification, so we expect our targets will span a large range of stellar types and colors.

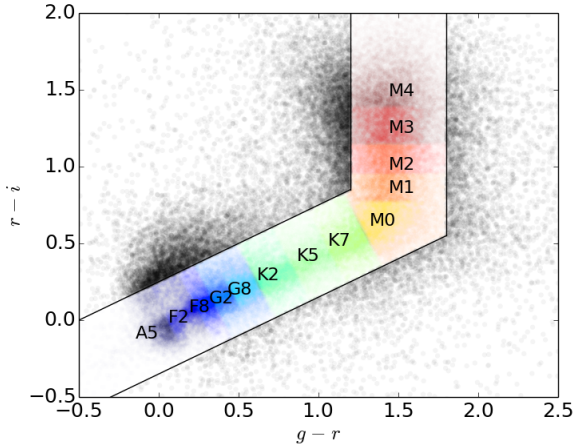


FIG. 16.— The SDSS $r-i$ versus $g-r$ distribution of all TDSS non-quasars (mostly stars). We approximate the main sequence and label and color-code different stellar types.

Fig. 16 shows the $r-i$ versus $g-r$ color distribution of all sources after removing the objects defined as quasars in Eq. 23. Statistically, we expect the vast majority of remaining objects to be stars. We match the objects to the SDSS main sequence from Kraus & Hillenbrand (2007) which we approximate as

$$\begin{aligned} r-i &= 0.5(g-r) - 0.05, \text{ for } g-r < 1.45, \\ r-i &> 0.675, \text{ for } g-r = 1.45. \end{aligned} \quad (25)$$

This is just a diagonal line which approximates the A through M0 stars and a vertical line that matches the colors of M1 and later stars. We classify our stars into categories defined in Table 6. These categories are chosen to be spaced at roughly 0.2 magnitude intervals in $g-r$, $r-i$ so that they are meaningful distinctions for a sample with error bars of just under 0.1 magnitudes. We set the location of the median subclass of star in Table 6 to the nearest point on the main sequence approximation in Eq. 25. We then match each star to the nearest stellar category median. The results are shown by the coloring

Stellar Class	Median Class	$g-r$	$r-i$	$(g-r)_{\text{line}}$	$(r-i)_{\text{line}}$
OBA	A5	-0.02	-0.17	-0.06	-0.08
Early F	F2	0.22	-0.01	0.19	0.05
Late F	F8	0.31	0.03	0.28	0.09
Early G	G2	0.42	0.11	0.40	0.15
Late G	G8	0.53	0.18	0.52	0.21
Early K	K2	0.71	0.29	0.70	0.30
Mid K	K5	0.95	0.44	0.96	0.43
Late K	K7	1.14	0.55	1.15	0.53
M0	M0	1.40	0.67	1.45	0.67
M1	M1	1.47	0.88	1.45	0.88
M2	M2	1.48	1.03	1.45	1.03
M3	M3	1.48	1.27	1.45	1.27
M4+	M4	1.48	1.51	1.45	1.51

TABLE 6

The different stellar categories shown in Fig. 16. We show the description, the median stellar subclass, the actual location of that subclass in $g-r$, $r-i$ space from Kraus & Hillenbrand (2007) and our approximation of this point on the main sequence approximation defined in Eq. 25.

in Fig. 16. We exclude stars that do not satisfy

$$\begin{aligned} r-i &> 0.5(g-r) - 0.35, \\ g-r &< 1.8, \\ r-i &< 0.5(g-r) + 0.25 \text{ or } r-i > 1.2 \end{aligned} \quad (26)$$

for tabulation purposes, these stars are called "Not MS" in Table 7 and are colored in greyscale in Fig. 16.

Table 7 lists the numbers and percentages of different stellar types shown in Fig. 16. It also presents the numbers and percentages of different stellar types after removing the "blue cloud" stars described in Eq. 25. These "blue cloud stars", if they are not actually quasars, are most likely F and G type stars.

There are two notable trends in Table 7. First, 17.1% of all objects are classified as "Not Main Sequence". This large fraction is perhaps not surprising since many of our variable targets will be interacting or eclipsing binaries, stars undergoing intense chromospheric activity or will otherwise have colors not consistent with simple stellar physics. Additionally, the fractions of variables are fairly constant across our stellar categories, ranging from 4.2% to 8.8%. There is no obvious reason for this to be the case. But it is convenient, as it will allow the study of a broad range of targets. Understanding why the fraction of stellar variables is constant in $r-i$, $g-r$ space will likely be a significant topic of interest for TDSS as spectra are analyzed.

Our stellar candidates are distributed much more uniformly across the main sequence than those presented in the Catalina Surveys Periodic Variable Star Catalog (Drake et al. 2014) and the analogous catalog from LINEAR (Palaversa et al. 2013). Specifically, a much larger fraction of our sources are redder K and M stars. The CSS and LINEAR teams require a period measurement for inclusion in their catalogs and are thus particularly sensitive to RR-Lyrae and other (mostly blue) pulsating variables with short periods. Since we do not require a period measurement, our sample includes many eclipsing binaries whose period is difficult to measure due to their low duty cycle. Eclipsing binaries occur across a wide range of stellar masses, so should be distributed rather uniformly across the main sequence. We also expect to find various flaring stars, especially towards the red end of the main sequence, which may not be periodic at all.

Stellar Class	N	P	$N_{\text{no bc}}$	$P_{\text{no bc}}$
OBA	4421	4.2	4406	4.5
Early F	6711	6.3	5240	5.4
Late F	6657	6.3	3951	4.0
Early G	6574	6.2	4006	4.1
Late G	6293	5.9	5507	5.6
Early K	6407	6.0	6405	6.5
Mid K	5014	4.7	5014	5.1
Late K	5857	5.5	5857	6.0
M0	8455	8.0	8455	8.6
M1	5894	5.5	5894	6.0
M2	7061	6.6	7061	7.2
M3	9380	8.8	9380	9.6
M4+	9390	8.8	9390	9.6
MS	88114	82.9	80566	82.4
Not MS	18190	17.1	17236	17.6
Previous SDSS Spectra				
Star	1742	1.6	1646	1.7
Galaxy	196	0.2	167	0.2

TABLE 7

The number and percentage of targets in the TDSS candidate list from different stellar classes/subclasses after removing all quasars (as defined by Eq. 23). N is the number of non-quasar targets of each type. P is the percentage of our total non-quasar targets from each stellar type. $N_{\text{no bc}}$ and $P_{\text{no bc}}$ are the analogous quantities for targets after objects in the "blue cloud" (Eq. 25) are also excluded. The first 13 rows add up to the main sequence (MS) line, and the total is of course 100%.

8.3. The Hypervariable Population

As mentioned in section 5.1, 1,108 of our sources are hypervariables with 2 or more magnitudes of variability, V (see Eq. 11). In Fig. 17, these variables have an unusual distribution of colors, with almost none near the quasar locus. These hypervariables are also significantly redder than our main population, suggesting that many of these stars may be cataclysmic variables, Mira variables or long-period variables.

We expect the hypervariables to be some of the most interesting objects in our survey and plan on examining this hypervariable population as well as the high variability stellar and quasar populations (mentioned as FES projects in the introduction). Specifically, we will examine the light curves from PS1 and shallower surveys like the Catalina Sky Survey, the Palomar Transient Factory and LINEAR (when available) and see how these relate to our early spectral identifications.

9. TDSS SELECTION FRACTION AS A FUNCTION OF COLOR

We can learn more about the TDSS selection algorithm by inverting the analysis in Section 8 and determining what percentage of objects with particular colors are selected as targets. Fig. 18 displays the selection percentage in the $g-r$, $u-g$ space from Fig. 11 and the $r-i$, $g-r$ space from Fig. 16. In this plot and in the accompanying tables below, we compare the total number of TDSS targets to the total number of objects in the TDSS footprint that pass our data quality cuts in Eq. 4. In broad strokes, the selection percentage is extremely low (0.3%) along the main sequence and much higher (above 10%) in areas of color space in which quasars or other more exotic astrophysical objects are expected to reside.

Table 8 tabulates the fraction of sources selected as variable objects in the categories in Fig. 18 (top) and Eq. 22. We only select 0.28% of objects on the main sequence, excluding the RR Lyrae box from which we select 0.61% of objects. Within the (very broad) quasar box (which includes many nonvariable, blue stars), we select 11.9%, although we select ap-

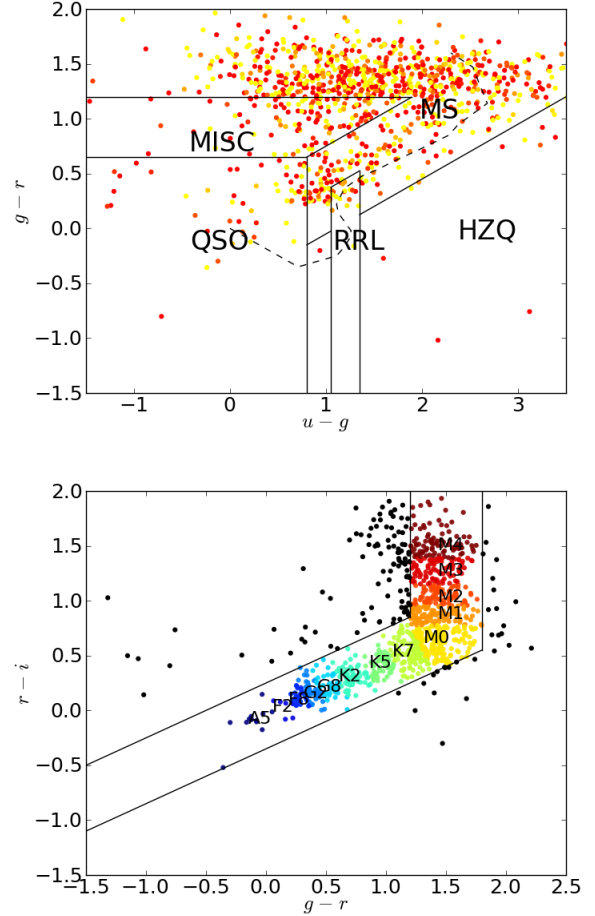


FIG. 17.— The SDSS $g-r$ versus $u-g$ distribution of all TDSS hypervariables (top) and the $r-i$ versus $g-r$ distribution of all TDSS hypervariables (bottom). In the top panel, low (high) priority objects are in red (yellow) and the QSO, MS, RRL and HZQ regions are the areas of color space that contain most quasars, main sequence stars, RR Lyrae stars and high-redshift quasars, respectively. In the bottom panel, we show the approximate positions of main sequence classifications.

Category	N_{targets}	$N_{\text{total objects}}$	% Selected
MS	75,754	27,079,176	0.28
QSO	143,052	1,201,995	11.90
RRL	7,358	1,204,246	0.61
HZQ	6,948	430,329	1.61
MISC	9,401	890,721	1.06

TABLE 8

Total number of targets, total number of objects and percentage-selected of different broad color-based categories as shown in Fig. 11 in our total TDSS sample.

proximately 30% of quasars with previous SDSS spectra as noted in Section 8. We select 1.61% and 1.06% of sources in the HZQ and MISC regions, respectively. These off-main sequence regions include variable subclasses like cataclysmic variables and white-dwarf main sequence binaries in addition to high-redshift quasars.

Table 9 tabulates the fraction of sources we select as variable objects from the categories in Fig. 18 (middle) and from section 8.2 after likely quasars are removed. We also present our results after removing the ambiguous "blue cloud" region from Eq. 25 in the right half of the table. Along the main

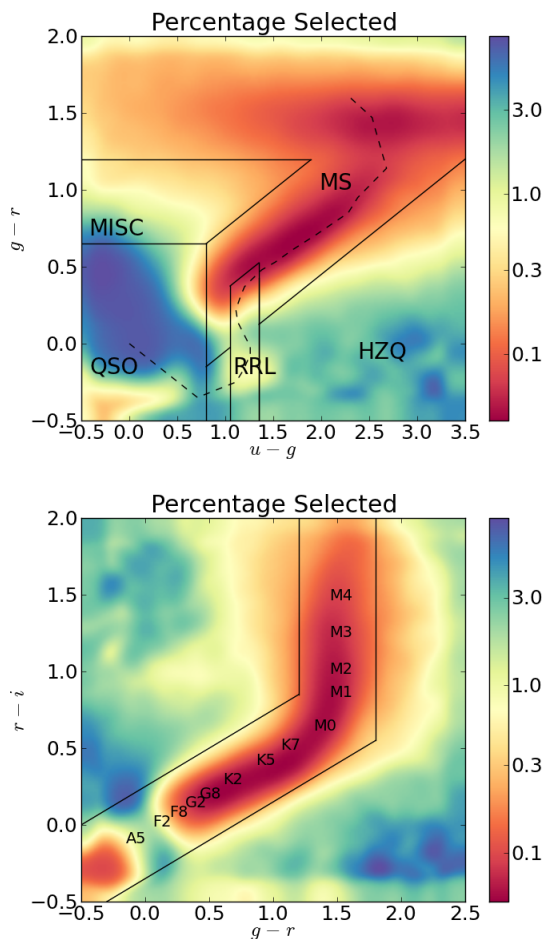


FIG. 18.— The percentage of SDSS objects which satisfy Eq. 4 that we select as TDSS targets as a function of $g-r$ and $u-g$ (top). The same percentage as a function of $r-i$ and $g-r$ (bottom). The variable object categories from Fig. 11 and the main sequence categories from Fig. 16 are also shown.

sequence, we preferentially select OBA stars (4.4%) and F stars (3.31%) over redder stars (0.2%-0.5%). Perhaps some of these early-type (blue) stars are the unusually colored quasars that remain after excluding our color-selected quasar sample, but the huge difference in selection percentage between early-type and late-type stars suggests that a relatively large fraction of early-type stars are early-type variables, including pulsators such as RR Lyrae stars.

10. STRIPE 82 AND CSS TARGETS WITH PREVIOUS SPECTROSCOPY OR VARIABILITY CLASSIFICATIONS

As a final probe into the TDSS sample, we run our algorithm on SDSS and PS1 data across the high Galactic latitude, $315^\circ < \text{RA} < 60^\circ$, area of SDSS Stripe 82 and cross-match our results with samples of objects with previous spectroscopy or variability classification. Both spectroscopy and known variable objects are significantly more dense in Stripe 82 than in the larger SDSS or eBOSS areas, so this dataset provides a relatively complete and homogeneous sample. We slightly modify our selection algorithm by using $2.5^\circ \times 2.5^\circ$ pixels with 62 TDSS-only targets per pixel since Stripe 82 is 2.5° wide. We then cross-match these targets (including shared CORE quasar targets) with $17 < i < 21$ point sources that have previous public SDSS spectroscopy and also cross-

match our sample with known variable objects. We use a set of 173 ellipsoidal/eclipsing binaries from Bhatti (2012), 235 RR Lyrae from Sesar et al. (2010) and 91 other low mass periodic sources from Becker et al. (2011). We also cross-match our complete target list with the union of the Catalina Sky Survey (CSS) periodic variables from Drake et al. (2014) and RR Lyrae variables from Drake et al. (2013) and Torrealba et al. (2015). This union contains 68,956 stellar variables, 5,978 of which satisfy the minimum data quality requirement from Eq. 4 and are in the TDSS area. Both our spectroscopic and variable object samples are the results of multiple different surveys with acute and intentional biases rather than a single statistically complete sample. The relative fractions of different sources that we detect are thus only suggestive of how our techniques will select various subclasses of variable objects.

Table 10 shows the numbers of objects of different spectroscopic types that pass our selection cut. We use SDSS spectroscopic pipeline classes (‘quasar’, ‘star’ or ‘galaxy’) and subclasses (of which there are many) rather than performing independent spectroscopic analysis. We combine all objects with spectroscopic type ‘quasar’ into the AGN category and classify them as either ‘AGN Broadline’ or ‘AGN Non-Broadline’. As expected, we select a significantly higher fraction of Broadline AGN. Many ‘Non-Broadline’ AGN are starburst galaxies or Seyfert type 2 galaxies in which the potentially variable central black hole is less dominant in the overall emission.

We only select 0.58% of objects with stellar spectra. This is also expected as most stars, unlike quasars, are not inherently variable. Conversely, only 358 of our approximately 2,400 stellar targets (15%) in Stripe 82 have previous spectra. The fact that 85% of our stellar targets are new, even in Stripe 82, an area with a disproportionately high density of spectra, emphasizes how large and unique the TDSS stellar sample is.

For convenience, we have bundled our stellar spectroscopic subclasses into the same photometric color subclasses we use in Tables 6 and 7 with additional categories for L and T dwarfs, carbon stars, cataclysmic variables and white dwarfs. Roughly half of the stars selected have OBA type colors. This population is highly weighted toward the ‘A’ end, and many of these stars are likely RR Lyrae or anomalous Cepheid variables. The list of stars with previous SDSS spectra is heavily biased towards OBA stars. Only 4.2% of our non-quasar targets are OBA targets. We also tend to select a relatively high percentage of L and T stars (3.24%) as well as carbon stars (3.42%), which are likely in binaries (Green 2013). We only select 3.16% of cataclysmic variables, objects that by definition have large variability amplitudes, but relatively short duty cycles. The L, T, carbon star and cataclysmic variable selection fractions are all suspect as a large number of objects are misidentified with these intrinsically rare classifications in the SDSS spectroscopic pipeline. In practice, objects identified by TDSS with these rare classifications may require additional observations to classify them with certainty. We select 4% of unresolved objects with galaxy spectra. These are probably intermediate AGN not recognized as quasars by the SDSS algorithm due to relatively weak emission lines, AGN with resolved galaxy flux that SDSS misclassified morphologically or occasionally supernova hosts.

Table 11 lists the fractions of previously identified Stripe 82 variable objects we detect. We only detect 15% of the Bhatti (2012) binaries. Binaries typically produce the ≈ 0.2 magnitudes of variability we require for targets only when they are nearly fully eclipsing and thus have a relatively low duty cy-

Stellar Class	N_{targets}	$N_{\text{total objects}}$	% Selected	$N_{\text{targets no bc}}$	$N_{\text{total objects no bc}}$	% Selected _{no bc}
OBA	4,421	100,368	4.40	4,406	100,219	4.40
Early F	6,711	202,730	3.31	5,240	88,923	5.89
Late F	6,657	1,407,242	0.47	3,951	455,731	0.87
Early G	6,574	2,408,184	0.27	4,006	1,479,662	0.27
Late G	6,293	2,810,731	0.22	5,507	2,668,135	0.21
Early K	6,407	2,755,323	0.23	6,405	2,755,301	0.23
Mid K	5,014	2,147,150	0.23	5,014	2,147,150	0.23
Late K	5,857	2,565,386	0.23	5,857	2,565,386	0.23
M0	8,455	3,465,678	0.24	8,455	3,465,678	0.24
M1	5,894	2,616,959	0.23	5,894	2,616,959	0.23
M2	7,061	2,944,265	0.24	7,061	2,944,265	0.24
M3	9,380	3,443,992	0.27	9,380	3,443,992	0.27
M4+	9,390	2,365,245	0.40	9,390	2,365,245	0.40
MS	88,114	29,233,253	0.30	80,566	2,7096,646	0.30
NMS	18,190	1,011,878	1.80	17,236	997,638	1.73
Previous SDSS Spectra						
Star	1,742	219,463	0.79	1,646	158,830	1.04
Galaxy	196	6,981	2.81	167	6,005	2.78

TABLE 9

The numbers and percentages of targets selected from different stellar classes/subclasses after removing all quasars (as defined by Eq. 23). N_{targets} is the number of non-quasar targets of each selected while N_{total} is the total number of non-quasar objects that pass our data quality requirements. The % Selected columns is the percentage of objects that we select in our total sample. We also show the analogous quantities for targets after objects in the "blue cloud" (Eq. 25) are also excluded (subscripted "no bc").

Spec Class	N_{S82}	ρ_{S82}	$N_{S82 \text{ TDSS}}$	$\rho_{S82 \text{ TDSS}}$	TDSS%
AGN	24,315	47.44	6,788	13.24	27.92
AGN Broadline	18,999	37.07	5,727	11.17	30.14
AGN Non-Broadline	5,316	10.37	1,061	2.07	19.96
Star	62,147	121.26	358	0.70	0.58
OBA	6,080	11.86	160	0.31	2.63
Early F	10,151	19.81	55	0.11	0.54
Late F	6,895	13.45	26	0.05	0.38
Early G	3,469	6.77	3	0.01	0.09
Late G	410	0.80	3	0.01	0.73
Early K	8,789	17.15	20	0.04	0.23
Mid K	432	0.84	3	0.01	0.69
Late K	3,177	6.20	4	0.01	0.13
M0	3,200	6.24	1	0.00	0.03
M1	2,696	5.26	10	0.02	0.37
M2	3,746	7.31	3	0.01	0.08
M3	4,485	8.75	8	0.02	0.18
M4+	6,274	12.24	28	0.05	0.45
L, T	556	1.08	18	0.04	3.24
Carbon Star	117	0.23	4	0.01	3.42
CV	253	0.49	8	0.02	3.16
WD	1,417	2.76	4	0.01	0.28
Galaxy	1,448	2.83	58	0.11	4.01

TABLE 10

A summary of SDSS spectroscopic pipeline classes and subclasses of all $315^\circ < \text{RA} < 60^\circ$, $17 < i < 21$ Stripe 82 point sources with spectroscopy. These columns are the number and density deg^{-2} of each type of object, the number and density deg^{-2} of each type of object that is selected by TDSS and the percentage of these objects that would be selected by TDSS. Many L, T, carbon star and CV classifications are suspect.

cle compared to the more constantly dynamic pulsators. More than half (51%) of the Sesar et al. (2010) RR Lyrae sample makes our cut. In fact, 156 of 235 (66%) of their RR Lyrae stars pass our ($E > 45.4$) RR Lyrae cut, with 15% being removed by our random downsampling in areas with more than 10 targets deg^{-2} . If the density of selected RR Lyrae stars here were applied over the whole sky, we would expect to find 1700 RR Lyrae stars. Additionally, our broad variability selector should identify many RR Lyrae stars whose light curves

Var Class	N_{S82}	ρ_{S82}	$N_{S82 \text{ TDSS}}$	$\rho_{S82 \text{ TDSS}}$	TDSS%
Binaries	173	0.34	26	0.05	15.03
RR Lyrae	235	0.46	120	0.23	51.06
Other Periodic	91	0.18	10	0.02	10.99

TABLE 11

The classes of selected $315^\circ < \text{RA} < 60^\circ$ Stripe 82 $17 < i < 21$ variable point sources. These columns are the number and density deg^{-2} of each type of object, the number and density deg^{-2} of each type of object that is selected by TDSS and the percentage of these objects that would be selected by TDSS.

Var Class	Num _{CSS}	Num _{CSS TDSS}	TDSS%
W-Ursae Majoris	1,982	550	27.75
Algol Eclipsing	364	47	12.91
β Lyrae	27	7	25.93
RR Lyrae	3,494	1,867	53.43
Blazhko	3	3	100.00
RS Canum Venaticorum	29	7	24.14
Anomalous Cepheid	3	2	66.67
Cepheid-II	11	3	27.27
High Amplitude δ Scuti	21	7	33.33
Long-Period Variables	7	3	42.86
Rotating Ellipsoidal	18	5	27.78
Post Common Envelope Binary	17	6	35.29
All	5,978	2,507	41.94

TABLE 12

The classes of selected periodic variable point sources from the Catalina Sky Survey. The columns are the number of each type of object in the TDSS area, the number detect by TDSS and the percentage of these objects that would be selected by TDSS.

are too faint to be precisely classified as RR Lyrae stars. It is likely that our estimate in Table 5 of 4,384 TDSS-only RR Lyrae targets made solely from photometry is not more than a factor of two too high. We only detect 11% of other periodic stars, likely due to their relatively small variability amplitudes.

We can perform a more in depth analysis for many of our sources over the full TDSS area by cross-matching with

known periodic variable objects from CSS. CSS is significantly shallower than PS1 (typical limiting magnitude of $V = 19.7$), and the CSS sources with measurable periodicity are biased towards the brighter end of the survey. Our sample of 5,978 CSS periodic variables analyzed by TDSS is heavily biased towards the bright end of the survey with 3,963 $i < 18$ and 5,621 $i < 19$ objects, respectively. Table 12 shows the numbers and percentages of CSS periodic variables selected by TDSS. The categories are those used by Drake et al. (2014). Our results here are similar to those in Stripe 82. In particular, we recover 53% of RR Lyrae and generally recover a large fraction of the pulsating stars (RR Lyrae, Blazhko stars, Cepheid variables, δ Scuti stars and Long period variables) which tend to have high amplitudes and duty cycles. As a reminder, we are randomly downsampling by 30%, so we should not exceed 70% completeness for a large population. We generally recover a smaller fraction of binary systems (W-Ursae Majoris, Algol Eclipsing, β Lyrae, RS Canum Venaticorum and Post Common Envelope Binaries) which tend to have lower duty cycles and amplitudes (although the categories here have relatively high amplitude).

As TDSS spectra are processed, we plan to compare our spectral identification of brighter TDSS-identified variable objects to those derived from higher cadence light curve analysis from other time domain imaging surveys (particularly the Catalina Sky Survey, the Palomar Transient Factory and LINEAR). Photometric classification of the stellar population may be supported through a machine-learning approach to the photometric time-series light curves. For example, the artificial neural-network based Eclipsing Binary Factory (EBF) pipeline (Paegert et al. 2014; Parvizi et al. 2014) has been used to automatically identify and sub-classify eclipsing binary stars in the *Kepler* field as eclipsing contact, eclipsing semi-detached, and eclipsing detached systems with a low false positive rate. These EBF sub-classifications are accompanied by a confidence level (i.e., posterior classification probability) for each target as a given variable type (e.g., Eclipsing Binary, Cepheid, δ Scuti, RR Lyrae). This EBF-generated confidence may then be used as quantitative corroboration for the spectral classification of TDSS stellar variable targets, and extrapolated cautiously to fainter targets.

11. CONCLUSIONS

TDSS promises to open a new window into the nature of astrophysical variable objects. Obtaining 220,000 $R \approx 2,000$, optical spectra will make TDSS a massive and unique spectroscopic survey of variable objects. Just as important as the scale of the TDSS sample is its breadth. By adopting a general variability metric and not selecting for specific types of variable objects in color space, TDSS will not only acquire spectra of 135,000 variable quasars, but it will also obtain spectra of 85,000 stellar targets including perhaps 4,000 RR Lyrae stars and 1,108 hypervariables (including blazars, CVs or other flaring stars), hundreds of carbon stars and multitudes of other variables yet to be determined. The TDSS stellar spectra have little overlap with previous SDSS stellar spectra and should prove to be a truly unique sample.

This survey is facilitated by the combination of SDSS and PS1 photometry. SDSS and PS1 both produce 10% level photometry out to $i = 21$ in the *griz* filters across an overlapping area of 14,400 deg², including the entire 7,500 deg² eBOSS area. The combination of an SDSS-PS1 photometry difference, spanning 6-10 years, and PS1-only variation, with time scales of hours to years, efficiently selects both long term vari-

able objects (quasars) and shorter term variable objects (most variable stars). After flagging and rejecting sources with unreliable photometry using sensible database queries, we use a Kernel Density Estimator and a Stripe 82 training set to produce a sample that we estimate to be 95% pure, based on Stripe 82 variability measurements. We suspect that our final sample will have even higher purity since some Stripe 82 non-variables may have simply been dormant during the epochs of Stripe 82 imaging but active during those of PS1. In addition, we increase purity further with visual image inspection. While the vast majority of our sample is selected in a relatively unbiased manner, we deliberately select 1,108 hypervariables (which vary by more than 2 magnitudes) and 73 *i*-dropouts to ensure that these potentially interesting objects are not excluded from our sample.

While precise and complete identification of variable objects is impossible with basic photometric colors, we analyze our sample in $u-g$, $g-r$, $r-i$ color space to characterize our sample in broad strokes. The majority of our sample (59%) resides in the traditional $z < 2.5$ quasar color region. However, after removing our overlap with the eBOSS CORE quasar sample and previous spectroscopy, only 13.4% of our TDSS-only targets reside in this region, while 76.1% of them lie along or near the main sequence (including 4.1% which are in the F-star region where most RR Lyrae lie). Our stellar population is spread out relatively evenly with 37.7% of our non-quasar sample being M stars, 40.9% being FGK stars, 4.2% being (intrinsically rare) OBA stars and 17.1% being outside our main sequence classifying scheme. This target diversity was a natural result of selecting objects based on their variability without explicit regard for their colors. Inverting this analysis, we select 11.9% of objects within a broad quasar color box while we only select 0.28% of main sequence stars. Within the main sequence, we select 4.4% of OBA stars, 3.31% of F stars and roughly 0.25% of all other stars.

We anticipate that the breadth of the TDSS sample will lead to a wide variety of applications. Our work here suggests variability will help improve quasar selection in redshift regimes where photometric color selection is difficult ($z \approx 2.8$) and distinguish white dwarfs from quasars. More interestingly, variability can help us identify quasars that are reddened by dust, have weakened emission lines or otherwise have unusual colors that mask them from conventional quasar searches. TDSS will also produce a relatively pure and complete quasar sample with respect to variability allowing a study of how quasar properties change with variability in a statistically robust way. Determining how the concentration of different types of stellar variables changes across the Milky Way will be a major survey goal of TDSS. TDSS also promises to produce the largest sample of outer Milky Way RR Lyrae spectra and will thus probe the outer halo with new precision. TDSS should also significantly expand our samples of cataclysmic variables and variable carbon stars, although confident identification may require additional observations, particularly for objects that are not in a quiet state when observed by TDSS. Finally, as the first truly large scale spectroscopic survey to access a broad range of variable types, TDSS serves as a pathfinder for future variability surveys like LSST, allowing both a statistical spectroscopic characterization of the variable object population and the identification of rare or extreme examples only found in large variable samples.

12. ACKNOWLEDGMENTS

Funding for the Sloan Digital Sky Survey IV has been provided by the Alfred P. Sloan Foundation and the Participating Institutions. SDSS-IV acknowledges support and resources from the Center for High-Performance Computing at the University of Utah. The SDSS web site is www.sdss.org.

SDSS-IV is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration including the Carnegie Institution for Science, Carnegie Mellon University, the Chilean Participation Group, Harvard-Smithsonian Center for Astrophysics, Instituto de Astrofísica de Canarias, The Johns Hopkins University, Kavli Institute for the Physics and Mathematics of the Universe (IPMU) / University of Tokyo, Lawrence Berkeley National Laboratory, Leibniz Institut für Astrophysik Potsdam (AIP), Max-Planck-Institut für Astrophysik (MPA Garching), Max-Planck-Institut für Extraterrestrische Physik (MPE), Max-Planck-Institut für Astronomie (MPIA Heidelberg), National Astronomical Observatory of China, New Mexico State University, New York University, The Ohio State University, Pennsylvania State University, Shanghai Astronomical Observatory, United Kingdom Participation Group, Universidad Nacional Autónoma de México, University of Arizona, University of Colorado Boulder, University of Portsmouth, University of Utah, University of Washington, University of Wis-

consin, Vanderbilt University, and Yale University.

The PS1 Surveys have been made possible through contributions of the Institute for Astronomy, the University of Hawaii, the Pan-STARRS Project Office, the Max-Planck Society, and its participating institutes, the Max Planck Institute for Astronomy, Heidelberg, and the Max Planck Institute for Extraterrestrial Physics, Garching, The Johns Hopkins University, Durham University, the University of Edinburgh, Queen's University Belfast, the Harvard-Smithsonian Center for Astrophysics, and the Las Cumbres Observatory Global Telescope Network, Incorporated, the National Central University of Taiwan, the National Aeronautics and Space Administration under Grant No. NNX08AR22G issued through the Planetary Science Division of the NASA Science Mission Directorate, the National Science Foundation under Grant No. AST-1238877, the University of Maryland, and Eotvos Lorand University (ELTE).

We thank Don York for many discussions, spanning a number of years, related to the combined SDSS and PS scientific potential, and we thank Tim Heckman for his support linking the two surveys through TDSS.

We gratefully acknowledge help with candidate visual inspection provided by Jerica Green (SAO) and Caroline Scott (Imperial).

REFERENCES

- Ahn, C. P., Alexandroff, R., Allende Prieto, C., et al. 2012, *ApJS*, 203, 21
- Aihara, H., Allende Prieto, C., An, D., et al. 2011, *ApJS*, 193, 29
- Akerlof, C., Amrose, S., Balsano, R., et al. 2000, *AJ*, 119, 1901
- Alcock, C., Allsman, R. A., Alves, D. R., et al. 2001, *ApJS*, 136, 439
- Antonucci, R. 1993, *% araa*, 31, 473
- Bañados, E., Venemans, B. P., Morganson, E., et al. 2014, *AJ*, 148, 14
- Badenes, C., Mullally, F., Thompson, S. E., & Lupton, R. H. 2009, *ApJ*, 707, 971
- Badenes, C., van Kerkwijk, M. H., Kilic, M., et al. 2013, *MNRAS*, 429, 3596
- Baliunas, S. L., Donahue, R. A., Soon, W. H., et al. 1995, *ApJ*, 438, 269
- Beck, P. G., Hambleton, K., Vos, J., et al. 2014, *A&A*, 564, A36
- Becker, A. C., Bochanski, J. J., Hawley, S. L., et al. 2011, *ApJ*, 731, 17
- Bhatti, W. A. 2012, PhD thesis, The Johns Hopkins University
- Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Science*, 327, 977
- Bovy, J., Myers, A. D., Hennawi, J. F., et al. 2012, *ApJ*, 749, 41
- Catelan, M., Minniti, D., Lucas, P. W., et al. 2011, in *RR Lyrae Stars, Metal-Poor Stars, and the Galaxy*, ed. A. McWilliam, 145
- Chambers, K. C. 2011, in *Bulletin of the American Astronomical Society*, Vol. 43, American Astronomical Society Meeting Abstracts 217, 222.02
- Comparat, J., Kneib, J.-P., Escoffier, S., et al. 2013, *MNRAS*, 428, 1498
- Connon Smith, R. 2007, *ArXiv Astrophysics e-prints*, astro-ph/0701654
- Dawson, K. S., Schlegel, D. J., Ahn, C. P., et al. 2013, *AJ*, 145, 10
- Debosscher, J., Blomme, J., Aerts, C., & De Ridder, J. 2011, *A&A*, 529, A89
- Deng, L.-C., Newberg, H. J., Liu, C., et al. 2012, *Research in Astronomy and Astrophysics*, 12, 735
- Drake, A. J., Djorgovski, S. G., Mahabal, A., et al. 2009, *ApJ*, 696, 870
- Drake, A. J., Catelan, M., Djorgovski, S. G., et al. 2013, *ApJ*, 763, 32
- Drake, A. J., Graham, M. J., Djorgovski, S. G., et al. 2014, *ApJS*, 213, 9
- Eisenstein, D. J., Weinberg, D. H., Agol, E., et al. 2011, *AJ*, 142, 72
- Fan, X., Narayanan, V. K., Lupton, R. H., et al. 2001, *AJ*, 122, 2833
- Fan, X., Strauss, M. A., Becker, R. H., et al. 2006, *AJ*, 132, 117
- Filiz Ak, N., Brandt, W. N., Hall, P. B., et al. 2013, *ApJ*, 777, 168
- Finkbeiner, D. P., Schlafly, E., & Green, G. 2014, in *American Astronomical Society Meeting Abstracts*, Vol. 223, American Astronomical Society Meeting Abstracts #223, 116.14
- Freedman, W. L., Madore, B. F., Gibson, B. K., et al. 2001, *ApJ*, 553, 47
- Friedman, J. H. 2001, *The Annals of Statistics*, 29, pp. 1189
- . 2002, *Computational Statistics & Data Analysis*, 38, 367
- Fukugita, M., Ichikawa, T., Gunn, J. E., et al. 1996, *AJ*, 111, 1748
- Geier, S., Hirsch, H., Tillich, A., et al. 2011, *A&A*, 530, A28
- Giveon, U., Maoz, D., Kaspi, S., Netzer, H., & Smith, P. S. 1999, *MNRAS*, 306, 637
- Green, P. 2013, *ApJ*, 765, 12
- Groot, P. J., Vreeswijk, P. M., Huber, M. E., et al. 2003, *MNRAS*, 339, 427
- Gunn, J. E., Carr, M., Rockosi, C., et al. 1998, *AJ*, 116, 3040
- Gunn, J. E., Siegmund, W. A., Mannery, E. J., et al. 2006, *AJ*, 131, 2332
- Hadjisjka, E., Rabinowitz, D., Baltay, C., et al. 2012, in *IAU Symposium*, Vol. 285, IAU Symposium, ed. E. Griffin, R. Hanisch, & R. Seaman, 324
- Hall, J. C., Henry, G. W., Lockwood, G. W., Skiff, B. A., & Saar, S. H. 2009, *AJ*, 138, 312
- Hubble, E. 1929, *Proceedings of the National Academy of Science*, 15, 168
- Ivezić, Ž., Smith, J. A., Miknaitis, G., et al. 2007, *AJ*, 134, 973
- Juric, M. 2011, in *Bulletin of the American Astronomical Society*, Vol. 43, American Astronomical Society Meeting Abstracts 217, 433.19
- Kaiser, N., Aussel, H., Burke, B. E., et al. 2002, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 4836, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, ed. J. A. Tyson & S. Wolff, 154
- Kaiser, N., Burgett, W., Chambers, K., et al. 2010, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7733, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series
- Kawaguchi, T., Mineshige, S., Umemura, M., & Turner, E. L. 1998, *ApJ*, 504, 671
- Knigge, C. 2011, in *Astronomical Society of the Pacific Conference Series*, Vol. 447, *Evolution of Compact Binaries*, ed. L. Schmidtobreick, M. R. Schreiber, & C. Tappert, 3
- Kordopatis, G., Gilmore, G., Steinmetz, M., et al. 2013, *AJ*, 146, 134
- Kraus, A. L., & Hillenbrand, L. A. 2007, *AJ*, 134, 2340
- Kunder, A., Koch, A., Rich, R. M., et al. 2012, *AJ*, 143, 57
- Law, N. M., Kulkarni, S. R., Dekany, R. G., et al. 2009, *PASP*, 121, 1395
- Lindgren, L., Babusiaux, C., Bailer-Jones, C., et al. 2008, in *IAU Symposium*, Vol. 248, IAU Symposium, ed. W. J. Jin, I. Platais, & M. A. C. Perryman, 217
- LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, *ArXiv e-prints*, arXiv:0912.0201 [astro-ph.IM]
- Lupton, R. H., Gunn, J. E., & Szalay, A. S. 1999, *AJ*, 118, 1406
- MacLeod, C. L., Ivezić, Ž., Kochanek, C. S., et al. 2010, *ApJ*, 721, 1014
- Mathur, S., García, R. A., Ballot, J., et al. 2014, *A&A*, 562, A124
- Morgan, C. W., Kochanek, C. S., Morgan, N. D., & Falco, E. E. 2010, *ApJ*, 712, 1129
- Morganson, E., De Rosa, G., Decarli, R., et al. 2012, *AJ*, 143, 142
- Morganson, E., Burgett, W. S., Chambers, K. C., et al. 2014, *ApJ*, 784, 92
- Mullally, F., Badenes, C., Thompson, S. E., & Lupton, R. 2009, *ApJ*, 707, L51
- Mumford, G. S. 1963, *AJ*, 68, 286
- Paegert, M., Stassun, K. G., & Burger, D. M. 2014, *AJ*, 148, 31
- Palanque-Delabrouille, N., Yèche, C., Myers, A. D., et al. 2011a, *A&A*, 530, A122
- . 2011b, *A&A*, 530, A122
- Palaversa, L., Ivezić, Ž., Eyer, L., et al. 2013, *AJ*, 146, 101
- Parvizi, M., Paegert, M., & Stassun, K. G. 2014, *ArXiv e-prints*, arXiv:1409.3237 [astro-ph.SR]
- Parzen, E. 1962, *Ann. Math. Statist.*, 33, 1065
- Pereyra, N. A., Vanden Berk, D. E., Turnshek, D. A., et al. 2006, *ApJ*, 642, 87
- Pojmanski, G. 2002, *Acta Astron.*, 52, 397
- Pourbaix, D., Knapp, G. R., Szkody, P., et al. 2005, *A&A*, 444, 643

- Pritchett, C. J., & van den Bergh, S. 1987, *ApJ*, 316, 517
- Rebassa-Mansergas, A., Nebot Gómez-Morán, A., Schreiber, M. R., Girven, J., & Gänsicke, B. T. 2011, *MNRAS*, 413, 1121
- Rees, M. J. 1984, *% araa*, 22, 471
- Richards, J. W., Starr, D. L., Butler, N. R., et al. 2011, *ApJ*, 733, 10
- Rodgers, A. W. 1957, *MNRAS*, 117, 85
- Rosenblatt, M. 1956, *Ann. Math. Statist.*, 27, 832
- Ruan, J. J., Anderson, S. F., Dexter, J., & Agol, E. 2014, *ApJ*, 783, 105
- Schatzman, E. 1962, *Annales d'Astrophysique*, 25, 18
- Schlaflly, E. F., Finkbeiner, D. P., Juric, M., et al. 2012, *ArXiv e-prints*, arXiv:1201.2208 [astro-ph.IM]
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *ApJ*, 500, 525
- Sesar, B., Ivezić, Ž., Lupton, R. H., et al. 2007, *AJ*, 134, 2236
- Sesar, B., Ivezić, Ž., Grammer, S. H., et al. 2010, *ApJ*, 708, 717
- Smee, S. A., Gunn, J. E., Uomoto, A., et al. 2013, *AJ*, 146, 32
- Smith, R. C. 1995, *Observational Astrophysics* (Cambridge University Press)
- Stephenson, C. B. 1960, *AJ*, 65, 60
- Tonry, J. L., Stubbs, C. W., Lykke, K. R., et al. 2012, *ApJ*, 750, 99
- Torreblanca, G., Catelan, M., Drake, A. J., et al. 2015, *MNRAS*, 446, 2251
- Udalski, A., Szymanski, M. K., Soszynski, I., & Poleski, R. 2008, *Acta Astron.*, 58, 69
- Ulrich, M.-H., Maraschi, L., & Urry, C. M. 1997, *ARA&A*, 35, 445
- Urry, C. M., & Padovani, P. 1995, *PASP*, 107, 803
- Vanden Berk, D. E., Wilhite, B. C., Kron, R. G., et al. 2004, *ApJ*, 601, 692
- Vivas, A. K., Zinn, R., Abad, C., et al. 2004, *AJ*, 127, 1158
- Wambsganss, J. 2006, in *Saas-Fee Advanced Course 33: Gravitational Lensing: Strong, Weak and Micro*, ed. G. Meylan, P. Jetzer, P. North, P. Schneider, C. S. Kochanek, & J. Wambsganss, 453
- Wang, X., Chen, X., Zheng, Z., et al. 2009, *MNRAS*, 394, 1775
- Wilson, O. C. 1963, *ApJ*, 138, 832
- Woźniak, P. R., Williams, S. J., Vestrand, W. T., & Gupta, V. 2004a, *AJ*, 128, 2965
- Woźniak, P. R., Vestrand, W. T., Akerlof, C. W., et al. 2004b, *AJ*, 127, 2436
- Wyrzykowski, L., Kostrzewa-Rutkowska, Z., Kozłowski, S., et al. 2014, *ArXiv e-prints*, arXiv:1409.1095
- Yanny, B., Rockosi, C., Newberg, H. J., et al. 2009, *AJ*, 137, 4377
- York, D. G., Adelman, J., Anderson, Jr., J. E., et al. 2000, *AJ*, 120, 1579
- Zasowski, G., Johnson, J. A., Frinchaboy, P. M., et al. 2013, *AJ*, 146, 81
- Zhevakin, S. A. 1959, *AZh*, 36, 996
- Zucker, D. B., de Silva, G., Freeman, K., Bland-Hawthorn, J., & Hermes Team. 2012, in *Astronomical Society of the Pacific Conference Series*, Vol. 458, *Galactic Archaeology: Near-Field Cosmology and the Formation of the Milky Way*, ed. W. Aoki, M. Ishigaki, T. Suda, T. Tsuchimoto, & N. Arimoto, 421

APPENDIX

COMPARISON OF KDE TO BOOSTED DECISION TREE

In order to investigate whether more complex techniques that utilize a greater variety of variability features can offer significant improvement over our variability-based KDE approach, we compared the KDE results with those obtained using a Stochastic Gradient Boosted Decision Tree (SGBDT) technique (Friedman 2001, 2002). Gradient boosting is one of the most powerful and commonly used machine learning techniques, and among its advantages are that it is highly flexible and fairly robust against overfitting. The basic idea behind gradient boosting is to build up a classifier (or regression function) as a linear combination of many weak classifiers. In most applications, including ours, the weak classifiers are shallow binary decision trees. One can think of the technique as modeling the logarithm of the probability that an object is a variable object, given the set of input variability features, as a basis expansion in a set of shallow decision trees, where each decision tree is derived sequentially from the training data. In the stochastic implementation that we used, the decision trees are derived sequentially using a random subsample of the training data, which improves the prediction error by reducing variance in the estimator through averaging. In addition to the median(*SDSS-PS1*), median(*Var*) and median(*mag*) features used in our standard selection algorithm, we add χ^2_{red} , Q_{tot} , ν and median(σ). Here, χ^2_{red} is the reduced χ^2 of our *PS1* $g_{\text{P1}}r_{\text{P1}}i_{\text{P1}}z_{\text{P1}}$ magnitudes assuming a constant for each of the $g_{\text{P1}}r_{\text{P1}}i_{\text{P1}}z_{\text{P1}}$ filters. Q_{tot} is the average of $Q_{75} - Q_{25}$ across *griz* filters, where Q_{75} and Q_{25} are, respectively, the 75th and 25th percentile *PS1* measurement in each filter. The quantity ν is a four filter white noise amplitude described in Morganson et al. (2014). Median(σ) is the median *PS1* standard deviation across the *griz* filters.

We used the stochastic gradient boosting algorithm implemented by the Python scikit-learn package⁰. There are a few tuning parameters in this algorithm. The first is the fraction of the training data that is used in each subsample when deriving each weak classifier. We set this parameter to 0.5, a recommended default value. Another tuning parameter is the learning rate, which controls the amount of shrinkage employed. A higher learning rate means that less shrinkage is applied to each of the base classifiers (shallow decision trees), and the model is built up faster. We adopt the default value of 0.1. The number of decision trees to use in the sum is chosen to be 84, found by minimizing the ‘out-of-bag’ error; the out-of-bag error is the error as evaluated by that subsample of the training set that was not used to build the next weak classifier. Finally, the maximum allowed depth of each decision tree in the sum was chosen to be 3, found to minimize the test error, where we withheld 25% of the Stripe 82 data set as test data and used the remaining 75% to train the algorithm. Ultimately, the SGBDT assigns every object in our $135^\circ < RA < 150^\circ$, $45^\circ < DEC < 60^\circ$ test set (as well as our training variable object and standard sets) a probability of being a variable object. This quantity is analogous to the *E* quantity (and related probability) defined in Section 4 for our KDE.

The SGBDT also provides a relative measure of the importance of each feature in classifying variable objects. The most important feature was found to be median($|SDSS-PS1|$), followed by median(*Var*) and median(σ). These three features contained approximately 60% of the total feature importance measure.

In Table 13, we show the SGBDT analog of Table 4. As in Section 4, we set thresholds in our SGBDT P_{var} so that 10, 20... 60 TDSS-only targets deg^{-2} pass the threshold in our test set. We can then count the number of variable objects and standards that pass these thresholds and calculate purities and other quantities with the same procedures described in Section 6. At the crucial density of 10 TDSS-only targets deg^{-2} (the density of our actual target list), the SGBDT sample is slightly more pure than our KDE sample (90.8% versus 86.4% in P_{tar}). However, our KDE performs significantly better at finding CORE quasars and objects with previous SDSS spectra and identifies 9.1 additional objects deg^{-2} . Since we are interested in the total sample that passes our threshold, this feature is a decisive advantage for the KDE. We also conceptually prefer using the KDE method which uses a few robust quantities that may be more homogeneous across our sample.

⁰ <http://scikit-learn.org>

$N_{\text{tar } 20}$	$N_{\text{tar test}}$	N_{QSO}	N_*	N_{lovar}	N_{CORE}	N_{prev}	N_{tot}	P_{tar}	P_{tot}
60	67.8	2.9	11.7	53.1	14.1	15.7	97.6	36.8	45.6
50	56.5	2.6	13.1	40.7	13.5	14.9	84.9	41.9	52.0
40	45.4	2.3	14.6	28.6	13.0	14.0	72.4	49.1	60.6
30	35.2	1.9	15.5	17.8	12.1	13.0	60.4	58.2	70.4
20	23.7	1.5	14.3	7.9	10.7	11.4	45.7	71.2	82.7
10	11.3	0.9	9.4	1.1	8.2	8.3	27.8	90.8	96.0

TABLE 13

The Stochastic Gradient Boosted Decision Tree analog of Table 4. Estimated target counts and purities from Stripe 82 tests at different variability cutoffs. All counts are in units of deg^{-2} . All purities are percentages. $N_{\text{tar } 20}$ is the number of targets in the 20th percentile pixel for a given threshold while $N_{\text{tar test}}$ is the number of targets in our test field. N_{QSO} , N_* and N_{lovar} are the estimated numbers of TDSS-unique quasars, stars and low-variability objects, respectively. N_{CORE} and N_{prev} are the estimated numbers of objects we share with the CORE quasar sample or have previous SDSS spectroscopy. N_{tot} is the total number of candidates. P_{tar} and P_{tot} are the estimated purities of our TDSS-only targets and our total targets, respectively.